Developmental disorders

Book chapter (Author's draft)

http://eprints.bbk.ac.uk/4609

Citation:

Karmiloff-Smith, A. and Thomas, M.S.C. (2002). 'Developmental disorders'. In Arbib, M.A. (Ed) *The Handbook of Brain Theory and Neural Networks* 2nd Edn. Cambridge, U.S.: MIT Press, pp.339-342

(ISBN: 9780262011976)

© 2002 Emerald

Publisher version

Developmental disorders

Annette Karmiloff-Smith

Neurocognitive Development Unit,

Institute of Child Health, London.

Michael S. C. Thomas

Neurocognitive Development Unit,

Institute of Child Health, London.

Running head: Developmental disorders

Address for correspondence:

Michael Thomas,

Neurocognitive Development Unit,

Institute of Child Health,

30, Guilford Street,

London WC1N 1EH, UK.

Email: M.Thomas@ich.ucl.ac.uk

Tel.: 0171 905 2747

Fax: 0171 242 7177

Introduction

Connectionist models have recently provided a concrete computational platform from which to explore how different initial constraints in the cognitive system can interact with an environment to generate the behaviors we find in normal development (Elman et al., 1996; Mareschal & Thomas, 2000). In this sense, networks embody several principles inherent to Piagetian theory, the major developmental theory of the twentieth century. By extension, these models provide the opportunity to explore how shifts in these initial constraints (or boundary conditions) can result in the emergence of the abnormal behaviors we find in atypical development. Although this field is very new, connectionist models have already been put forward to explain disordered language development in Specific Language Impairment (Hoeffner & McClelland, 1993), Williams Syndrome (Thomas & Karmiloff-Smith, 1999), and developmental dyslexia (Seidenberg and colleagues, see e.g. Harm & Seidenberg, in press); to explain unusual characteristics of perceptual discrimination in autism (Cohen, 1994; Gustafsson, 1997); and to explore the emergence of disordered cortical feature maps using a neurobiologically constrained model (Oliver, Johnson, Karmiloff-Smith, & Pennington, in press). In this entry, we will examine the types of initial constraints that connectionist modelers typically build in to their models, and how variations in these constraints have been proposed as possible accounts of the causes of particular developmental disorders. In particular, we will examine the claim that these constraints are candidates for what will constitute innate knowledge. First, however, we need to consider a current debate concerning whether developmental disorders are a useful tool to explore the (possibly innate) structure of the

normal cognitive system. We will find that connectionist approaches are much more consistent with one side of this debate than the other.

Developmental disorders and modularity

Cognitive neuropsychology assumes that the adult cognitive system has a modular structure, and that selective behavioral deficits in adults with brain damage can reveal this modular structure. Developmental disorders can produce apparently specific deficits in the endstate of development. For example, Williams syndrome (WS), a developmental disorder caused by a microdeletion of contiguous genes on one of the alleles of chromosome 7, is characterized by a behavioral profile of relative proficiency in language, face processing, and theory-of-mind (attributing mental states to others), but severe deficits in other skills such as visuospatial processing, number, and problem solving (Karmiloff-Smith, 1998). In hydrocephalus with associated myelomeningocele (a protrusion of the membranes of the brain or spinal cord through a defect in the skull or spinal column), language can be the only area of proficiency. Subjects suffering from Specific Language Impairment (SLI) show the opposite pattern, performing within the normal range in all domains except language. In autism, even individuals with normal IQs are selectively impaired in tasks that require judging another's mental states (Baron-Cohen, Tager-Flusberg, & Cohen, 1993).

The specific deficits in the endstate of development and the genetic origin of many of these disorders have encouraged some researchers to use developmental disorders as if they were cases of adult brain damage, and as if the particular deficits might reveal the innate modular structure of the cognitive system. For example, Baron-

Cohen et al. (1993) have argued that, in individuals with autism, an apparent deficit in reasoning about mental states can be explained by the impairment of an innate, dedicated module for such reasoning (the 'Theory of Mind' module). Van der Lely (1997) maintains that behavioral deficits in the language performance of children with grammatical SLI can be explained by damage to an underlying, innate module representing syntactic (rule-based) information. Clahsen and Almazan (1998) have proposed that a behavioral deficit in WS language supports the view that while their syntactic skills are 'intact', they have a deficit to a specific aspect of their (modular) language knowledge, that of accessing information about words that are exceptions to syntactic rules.

However, there are a number of problems with the adult brain damage approach to developmental disorders (Karmiloff-Smith, 1998). These boil down to the suspicion that such an approach massively underestimates the complexity of the path from gene to behavior. According to current genetic knowledge, there are no area-specific genes involved in the construction of specific neocortical areas that might serve the function of coding directly for domain-specific developmental outcomes. Even if there were, the idea that behavioral deficits identified in the endstate of a developmental disorder could reflect the impairment of a single module is predicated on the assumption that the rest of the cognitive system could nevertheless develop normally. For this to be true requires either that modules develop independently of overall brain growth, or that the content of modules is fixed in advance (i.e. the content is innately specified). But neither of these assumptions is likely to be true. With regard to the first, Bishop (1997) has argued persuasively that interactivity rather than independence is the hallmark of <u>early</u> develop-

ment. With regard to the second, it seems likely that modular structure in the cognitive

system and in the brain is an outcome of development rather than a precursor to it, and

that the neonate brain does not support innate representations with specific content

(Elman et al., 1996; Johnson, 1997). A growing number of studies show how both neural

localization and neural specialization for biologically important functions, such as species

recognition (Johnson, 1997) and language (Neville, 1991), take place gradually across

development.

The alternative to the use of the adult brain damage model is the neuroconstruc-

tivist approach (Elman et al., 1996; Karmiloff-Smith, 1998) which views developmental

disorders in terms of different developmental trajectories, caused by initial differences at

a neurocomputational level. Thus there might be differences in the microcircuitry of the

brain or the firing properties of neurons, as opposed to discrete lesions to particular large-

scale brain structures or pathways. In this view, development is an interactive process in

which the cognitive system self-organizes in response to interactions with a structured

environment. Interestingly, this approach suggests that people with developmental

disorders may exhibit strengths as well as weaknesses. This prediction is consistent with

superior face recognition skills found in WS and superior perceptual discrimination

abilities found in autism. Neuroconstructivism further suggests that equivalent behavior

across normal and abnormal phenotypes may mask different underlying cognitive proc-

esses. The notion that an ability is 'intact' or 'spared' because there is no apparent deficit

at the behavioral level employs terminology from the adult brain damage model that may

be misleading. To take an example, people with Williams syndrome can display scores

on some language and face processing tasks which are in the normal range. Nevertheless,

closer examination suggests different <u>cognitive</u> processes underlie the equivalent <u>behav-ioral</u> scores (Karmiloff-Smith, 1998).

Initial constraints in connectionist models

Current connectionist models of developmental disorders employ alterations to the initial model constraints which, after training, lead to an endstate exhibiting behavioral deficits. Connection weights are usually randomized so that the normal network has no specific knowledge prior to training. Thus it follows that the atypical network has <u>no specific deficit in knowledge</u> prior to training. The behavioral deficits that emerge when these atypical networks are trained are often quite different to the effects of damaging a normal network model after training has been completed. This holds even when the network manipulations are the same in each case. Thus current connectionist models are more consistent with the neuroconstructivist approach to developmental disorders than the adult brain damage approach.

What, then, are the initial constraints that connectionist modelers build into their models of cognitive development? And how do changes to these constraints alter the trajectory of development? In fact, the constraints that connectionist models build in are quite strong ones, and this may come as a surprise to some readers. Connectionist models are often mis-characterised as being unitary / homogeneous / seamless / undifferentiated, domain-general learning devices, whereby the environment is all-powerful in shaping the behavior of the final system. In fact, current connectionist models have a great deal of pre-existing structure built into them prior to any exposure to their training environment. What is <u>general</u> about connectionism are the principles of computation (Seidenberg,

1993). When the general principles of computation are combined with the boundary

conditions for a specific domain, the result is a domain-specific model. It is the generality

of these principles that gives the connectionist approach its explanatory power. That is,

connectionism seeks not just to formulate descriptive generalizations about empirical

phenomena, but to show how they derive from underlying and independently motivated

principles (Seidenberg, 1993). However, connectionist models are just as reliant on the

constraints of a given domain as they are on the computational principles. Without

justified limitations in the design of network models, they become overpowerful data-

fitting devices that can at best provide descriptively adequate accounts of cognitive

abilities. In short, connectionist models of development do include initial structure but

not initial representational content. The point is that in interacting with a training

environment, networks create representational content and become <u>increasingly</u>

structured. This additional structure reflects the nature of the training environment.

      The structure or boundary conditions that these models build in prior to training

typically involve the following:

(1) <u>The initial state of the network, in terms of the number of units, layers, connections,</u>

    <u>and the pattern of connectivity, collectively known as the network architecture</u>. The

    architecture determines the computational power of the network, and the type of

    computations to which the network will be suited. For example, recurrent networks

    are suited to processing sequential information, whereas associative networks are

    suited to pattern recognition. The <u>a priori</u> choice of the architecture will have a central

    role in determining the adequacy of the network in modeling a given domain of

cognitive development. A reduction in the number of processing units, or the elimination of internal (hidden) processing units, will restrict the computational power of the network and, depending on the nature of the domain, cause some or all parts of the problem to be learned inadequately. Addition of layers of internal units beyond a single layer tends to delay learning, without a marked increase in effective computational power. Increase of hidden units within a single layer tends to improve performance on the training set, but impair generalization beyond the training set.

(2) <u>The way a particular cognitive problem is presented to the network, in terms of the input and output representations</u>. For a given domain, the representations determine the nature of the computational problem that the network will face. When the network has to extract a function from the training set (such as a general rule), the representational scheme will be crucial in determining how transparent or opaque this function is to the network. For instance, if a network is given a training set in a form that masks the relevant similarity between those items in the problem domain that obey a rule, the network will have difficulty in extracting this rule.

(3) <u>The learning algorithm that the network will use to change its connection weights (and potentially, its architecture)</u>. Most networks are trained by changing weights to minimize some cost function, such as the difference between the actual output and a target output. The rate at which weights are changed can have an impact on the success of a network in learning a problem. In particular, in complex domains, if weights are changed too quickly, the network may commit too early to a partial

solution to the problem, and be resistant to change with subsequent training. The learning algorithm is key in determining the plasticity of the network to further learning. Some algorithms allow on-line changes to network architecture depending on how well the network is learning a problem. The way that the network's computational power is altered on-line will again have a considerable influence on the success of the network in capturing a cognitive ability (see (1) above).

(4) <u>The regime of training which the network will undergo</u>. After modelers determine the network and the learning rule, they then expose the network to a training set. Often the network is exposed to the entire training corpus from the start of training. However, in some cases, the network might be trained on an initially limited training set, perhaps based on assumptions about the nature of a child's early learning environment. This initial restriction will affect later training. It may aid learning if the smaller set is representative of the larger set, or if it allows the construction of internal representations that will be useful in learning the larger set. On the other hand, it may impair learning, if the initial training set contains detail irrelevant to the full domain. Alterations in network parameters early in training may have the same effect as restricting the initial training set, as Elman's work on learning syntax with recurrent networks has demonstrated (Elman et al., 1996).

In the connectionist framework, these constraints represent some of the candidates for innateness. Alterations in one or more of these constraints may then lead to the emergence of disordered representations and impaired behavior in a model of atypical

development.

We have noted that these models do not support innate representational content in that their weights are initially randomized. However, it is an open question about whether computational constraints (along with sensory input determined by the individual's interaction with the environment) are sufficient to drive development. One possible addition is the postulation of innate attentional predispositions. In this theory, innate knowledge is built into the subcortical part of brain in the form of a low acuity predisposition to attend certain stimuli. This predisposition then guides input that will subsequently drive learning in the more powerful cortical system. For example, such an innate predisposition in face recognition might encourage the newborn infant to attend preferentially to visual stimuli containing a single blob positioned centrally below two blobs (see Johnson, 1997). Innate predispositions provide another candidate factor that might be altered in the start state of the atypical system.

In general, current connectionist models of normal development do not restrict themselves to computational constraints and innate attentional predispositions in their start states, since these models incorporate high-level, domain-specific representations. These models must therefore be seen as a halfway house. In future they must be extended to show how their domain-specific representations may emerge from some prior process operating over lower level information (and with its own computational constraints).

We now turn to consider recent examples of connectionist models of developmental disorders.

Recent models

Autism. Cohen (1994) and Gustafsson (1997) have put forward connectionist accounts of associated perceptual differences in autism. These include heightened perceptual discrimination, over-attention to surface features, and a failure to extract deep structure. Cohen used a feedforward backpropagation network trained on a classification task, and varied the number of hidden units in the start state. With 'too few' hidden units, learning was poor. With 'too many' hidden units, the network learned quickly on the training set, but focused on idiosyncratic features of that set and was thus unable to generalize well to novel patterns. Cohen suggested that this pattern fits with the deficits found in different domains with children with autism. Notably, he reviewed post-mortem neuroanatomical evidence that suggested increased and decreased levels of neuronal density in the brains of individuals with autism, offering some justification for his manipulation. Gustafsson (1997) explained the combination of the failure to generalize with associated heightened perceptual discrimination as an impairment in the development of cortical feature maps. Although he ran no simulations, he suggested that higher than normal levels of within-layer inhibition in the initial cortical structure would lead to overly fine-tuned perceptual features, which would allow good discrimination but be unable to support good generalization. It is not clear, however, how either of these approaches may be readily extended to account for the high-level deficits found in autism.

Self-organizing feature maps in neocortex. Oliver et al. (in press) recently examined the ways in which just such a process of feature map formation could be disrupted by changes in the initial properties of a self-organizing connectionist network. They em-

ployed a neurobiologically constrained network in which a two-dimensional output layer received information from a single input retina. The network was shown a set of stimuli in the form of bars lying across the input retina. Oliver et al. showed that, using their initial parameter set, the output layer formed a topographic map of the possible inputs: certain areas of the output layer specialized in responding to each input, and areas representing similar inputs were adjacent to each other in the output layer. Oliver et al. then re-ran the model, disrupting the network in different ways prior to exposing it to the training stimuli. They varied the threshold of the output units, disrupted the connectivity between the input and output layers, disrupted the connectivity responsible for lateral inhibition in output layer, and changed the similarity of the input stimuli to each other. Importantly, these manipulations demonstrated that tiny differences in the initial constraints under which the model developed could have a very significant impact on the outcome of development. The resulting topographic map suffered a range of disruptions, including output units failing to specialize at all or simply turning off, specialization emerging but not in organized areas, and organized areas emerging but without adjacent areas representing similar looking bars.

Developmental dyslexia. This disorder has been the focus of much connectionist research, given the success of models in capturing the normal processes of reading. A number of models by Seidenberg and colleagues have sought to change initial constraints in reading models to simulate either surface dyslexia (where the subject has difficulty reading words which are exceptions to normal rules of pronunciation), phonological dyslexia (where the reading of novel words is impaired), or a combination of both types. Typically these models learn mappings between codes representing orthography,

phonology, and semantics. Surface dyslexia has been simulated by employing 'too few'

hidden units in the model. Phonological dyslexia has been simulated by degrading the

phonological representations in some way, for instance in the type of coding scheme

used. For example, Harm and Seidenberg (in press) pre-trained one part of their model to

develop appropriate phonological representations, prior to learning the reading task.

When this 'phonological' part of the model was impaired, either by reducing its initial

computational power or by limiting the size of the connection weights it could develop,

the result was a network exhibiting phonological dyslexia at the end of training on the

reading task.

Specific Language Impairment. Hoeffner and McClelland (1993) sought to capture

deficits found in the morphosyntax of subjects with SLI, specifically their difficulty in the

learning of rule-based inflectional morphology in verbs. Hoeffner and McClelland

employed an attractor network mapping between semantic codes and phonological codes.

They simulated SLI by changing the initial phonological representations, in line with a

hypothesis that SLI may be caused by early perceptual impairments. Specifically, they

impaired the network's ability to represent word-final stops and fricatives (including /t/,

/d/, and /s/). Although the model they used didn't show an ideal fit to the normal data

when unimpaired, it nevertheless captured a number of the key deficits of SLI when

trained with impaired representations, particularly a selective difficulty with the

formation of regular (+ed) past tenses. In this case, the initial phonological deficit

obscured precisely the information that the network required to be able to learn the

relevant generalizations about regular past tense formation.

<u>Williams Syndrome.</u> Recent work in our laboratory has examined underlying deficits in the language of individuals with Williams Syndrome. Initially thought to be 'spared', closer examination revealed a number of subtle deficits in their language. It had been reported that they show difficulties in forming the past tense of irregular verbs, while showing good performance on the regular, rule-based past tense formations (Clahsen & Almazan, 1998). Although our recent empirical work suggests that much of this apparently selective deficit is due to an overall delay in language development (young children also find irregular verbs difficult), we have recently used connectionist models to explore the hypothesis that changes in the nature of WS phonological representations may be one cause of the delay in their language development, and also cause an exaggerated difficulty with irregular verbs (Thomas & Karmiloff-Smith, 1999). In particular, in line with evidence of a relative strength in auditory short-term memory as young as 2½ years (the earliest age tested) and the reported hypersensitivity of the auditory system in adults with Williams syndrome, we altered the initial phonological representations in a 'normal' model of the development of past tense formation. As a result, there was a reduction in the similarity between the sounds making up each word. The model was trained to map between the phonological form of the present and past tense of verbs. We found that changes to the representational scheme prior to training produced both the delayed development and a selective deficit for irregular verbs, a behavioral pattern reported in WS. In this case, an initial difference in phonological representations differentially impaired the model's ability to learn two types of computations. Both regular and irregular past tenses took longer to learn, but in the

endstate, only irregulars showed a persisting deficit.

Developmental double dissociations

All of the models described here altered boundary conditions prior to training and, as a result, produced a different system in the endstate. The final performance of these models did indeed show 'selective' behavioral deficits. For example, it is a robust finding that in a system required to learn both regular and exception mappings, a limitation on internal system resources will tend to cause a differential impairment in the performance on exception mappings at the end of training. Such a system in the endstate is not a 'normal' system with the module responsible for learning exceptions somehow damaged or missing. Notably, the connectionist models of SLI and WS suggested that two different initial alterations to the phonological representations employed in past tense models can produce a form of double dissociation between regular and irregular past tense formation in the endstate. Temple (1997) has argued that such developmental double dissociations are key to identifying innate fractionations in the cognitive system. Following this logic, we might assume the existence of a selective deficit to a 'rule mechanism' in one syndrome and to an 'exception mechanism' in the other. Yet neither connectionist model employed a separate mechanism for each verb type. Once more, connectionist models suggest that we must rethink what counts as a valid inference from data to theory (Seidenberg, 1993).

Conclusion

Developmental disorders can inform the study of normal development because they

provide a broader view of the parameter space within which development takes place. The empiricist might claim that the environment was such a strong specifier of a capacity that systems with a wide variation of initial structures must come to reflect it. The nativist might claim that the environment is such a poor specifier of a capacity that the system must be pre-specified beforehand if it is to find the correct solution given the environmental input. However, the neuroconstructivist would argue that the robustness of the cognitive system to changes in its initial set-up (as long as we can come to understand precisely what these changes are) will reveal much about how evolution has placed its bets with regard to those capacities that can be trusted to emerge through experience, and those capacities which must be given a firmer guiding hand through development. Connectionist models provide a powerful tool with which to investigate the role of initial computational constraints in determining the trajectory of both typical and atypical development.

References

Baron-Cohen, S., H. Tager-Flusberg, and D. J. Cohen, 1993. <u>Understanding Other Minds: Perspectives from Autism,</u> Oxford: Oxford University Press.

Bishop, D. V. M. 1997. Cognitive neuropsychology and developmental disorders: Uncomfortable bedfellows, <u>Quarterly Journal of Experimental Psychology</u>, 50A:899-923.

Clahsen, H. and M. Almazan, 1998. Syntax and morphology in Williams Syndrome, <u>Cognition</u>, 68:167-198.

Cohen, I. L. 1994. An artificial neural network analogue of learning in autism, <u>Biological Psychiatry</u>, 36:5-20.

* Elman, J. L., E. A. Bates, M. H. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett, 1996. <u>Rethinking innateness: A connectionist perspective on development</u>, MIT Press, Cambridge (Mass).

Gustafsson, L. 1997. Inadequate cortical feature maps: A neural circuit theory of autism, <u>Biological Psychiatry</u>, 42:1138-1147.

Harm, M. W. and M. S. Seidenberg, in press. Phonology, reading acquisition, and dyslexia: Insights from connectionist models, <u>Psychological Review</u>.

Hoeffner, J. H. and J. L. McClelland, 1993. Can a perceptual processing deficit explain the impairment of inflectional morphology in developmental dysphasia? A computational investigation, Proceedings of the 25th Child language research forum, (E. V. Clark, ed.), Stanford University Press.

* Johnson, M. H. 1997. Developmental Cognitive Neuroscience, Blackwells, Oxford, UK.

* Karmiloff-Smith, A. 1998. Development itself is the key to understanding developmental disorders, Trends in Cognitive Sciences, 2:389-398.

* Mareschal, D. and M. S. C. Thomas, 2000. Self-organization in normal and abnormal cognitive development, in Brain and Behavior in Human Development. A source book, (A. F. Kalverboer and A. Gramsbergen, eds.), Dordecht: Kluwer Academic Publishers.

Neville, H. J. 1991. Neurobiology of cognitive and language processing: Effects of early experience, in Brain maturation and cognitive development: Comparative and cross-cultural perspectives. Foundation of human behavior, (K. R. Gibson and A. C. Peterson, eds.), New York: Aldine de Gruyter, pp. 355-380.

Oliver, A., M. H. Johnson, A. Karmiloff-Smith, and B. Pennington, in press. Deviations in the emergence of representations: A neuroconstructivist framework for analyzing developmental disorders, <u>Developmental Science</u>.

Seidenberg, M. 1993. Connectionist models and cognitive theory, <u>Psychological Science</u>, 4:228-235.

* Temple, C. 1997. <u>Developmental cognitive neuropsychology</u>, Psychology Press.

Thomas, M. S. C. and A. Karmiloff-Smith, 1999. Connectionist modeling of past tense formation in Williams syndrome, <u>Presentation to the Cognitive Psychology Section of the British Psychological Society</u>, York, UK, September 1999.

Van der Lely, H. K. J. 1997. Language and cognitive development in a grammatical SLI boy: Modularity and innateness, <u>Journal of Neurolinguistics</u>, 10:75-107.