Draft chapter to appear in: Mareschal, D., Johnson, M., Sirios, S., Spratling, M., Thomas, M. S. C., & Westermann, G. (forthcoming). Neuroconstructivism: How the brain constructs cognition. Oxford: Oxford University Press.

Neuroconstructivist principles applied to atypical development

(MT 10/7/03, revised 22/12/03)

In this chapter, we consider the implications of the neuroconstructivist principles for the study of development disorders.

Developmental disorders can be classified into four groups: genetic disorders caused by well understood genetic abnormalities (e.g., Fragile X syndrome, Down's syndrome, Williams syndrome, Turner's syndrome); disorders defined by one or more behavioural deficits (e.g., developmental dyslexia, Specific Language Impairment, autism); mental retardation of unknown aetiology; and disorders resulting from environmental factors (e.g., an impoverished environment, Foetal Alcohol syndrome). The first and last of these groups distinguish the locus of initial causality in terms of a nature/nurture distinction. The middle two groups tell us about the level of our current understanding of such disorders. For example, disorders like Specific Language Impairment (SLI) and autism appear to have a genetic component but the genes involved have not yet been identified (Bishop, North, & Donlan, 1995; Pennington & Smith, 1997; Simonoff, Bolton, & Rutter, 1998).

The study of developmental disorders proceeds with two aims in mind. The first of these is to identify appropriate methods of remediation and for behaviourally

defined disorders, early diagnosis to maximise the impact of remediation programmes. The second aim is to use disorders to help our understanding of the normal processes of development. We will shortly argue that developmental disorders are best conceived of as development occurring under atypical neurocomputational constraints. A successful research programme could use developmental disorders to throw into relief the form and potential variability of these constraints. Within this framework, "normal" development would simply constitute a special case of the settings of the constraints that guide all processes of development, successful or otherwise.

In relation to the second aim, little progress has been made in understanding the cognitive basis of general learning disability (mental retardation) where performance is lowered across all cognitive domains, let alone the neural bases underpinning such lowered performance. Disorders that show an uneven cognitive profile in their endstate offer most promise of theoretical insights. A number of disorders demonstrate dissociations in behaviour across different cognitive domains in adulthood. For example, Williams syndrome (WS) is characterised by a behavioural profile of relative proficiency in language and face processing, 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 (Karmiloff-Smith, 1998). Individuals suffering from 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). In Fragile X syndrome (FraX), the adult cognitive profile is characterised by relative strengths in vocabulary, long-term memory and holistic information processing but relative weaknesses in visuospatial cognition, attention, short-term memory and sequential information processing (Cornish, Munir, & Cross, 1999, 2001; Freund & Reiss, 1991).

Some genetic disorders are caused by fairly circumscribed genetic mutations. For instance WS is caused by a microdeletion of approximately 25 genes from one copy of chromosome 7 (Frangiskakis et al., 1996; Donnai & Karmiloff-Smith, 2000; Tassabehji et al., 1996, 1999). FraX is caused by the duplication of genetic material (the CGG repeat) in the 'Fragile X Mental retardation' (FMR1) gene on the X chromosome, preventing the reading of the DNA message that the gene encodes (O'Donnell & Warren, 2002). The absence of this gene's product is the sole genetic cause of the disorder. The combination of circumscribed genetic causes and uneven cognitive profiles mean that these disorders have the potential to illuminate links between genotype and phenotype. However, the correct explanatory framework for this endeavour has been a matter of some debate.

Developmental versus maturational accounts of uneven cognitive profiles

In the chapter on embrainment, we discussed the maturational perspective on functional brain development, in which newly emerging sensory, motor, and cognitive functions in the developing child are related to the independent maturation of areas of the brain (usually cerebral cortex) responsible for each function. We argued that the perspective is limited because the emergence of new behavioural skills is associated with widespread changes across many regions of the cortex, and functional brain

development appears to involve both increasing specialisation and localisation. Nevertheless, the maturational viewpoint has been a popular one within which to conceive of developmental deficits. Selective cognitive deficits are viewed as isolated failures of particular functional modules to development. For example, Baron-Cohen et al. (1993) have argued that in individuals with autism, an apparent deficit in reasoning about mental states can be explain by the impairment of an innate, dedicated module for such reasoning (the 'Theory of Mind' module). Van der Lely (1997) maintains that behavioural deficits in the language performance of children with so-called grammatical SLI can be explained by damage to a genetically determined, specialised module for processing syntactic (rule-based) information. Clahsen and Almazan (1998) have proposed that in WS language, while syntactic skills are considered to develop normally, there is a deficit in a component of the modular language system involved in accessing information about words that are exceptions to syntactic rules.

This conception of developmental deficits in effect seeks to extend the explanatory framework of adult cognitive neuropsychology to the developmental realm. Patterns of deficits in adults with brain damaged are interpreted in terms of intact and impaired functional modules. This framework and its methods are powerful tools for exploring cognitive deficits at a given point in time (Jackson & Coltheart, 2001). However, because the framework deals in static snapshots, its power to evaluate the distal origins of deficits is limited, and in the case of developmental deficits, use of the framework leads to the curious postulation of explanations that exclude the process of development (see discussion in Karmiloff-Smith, 1998; Thomas & Karmiloff-Smith, 2002a).

Somewhat more problematically, the static framework sometimes makes assumptions about development that are quite unlikely. For example, if interactive specialisation is the appropriate view of functional brain development, then it appears implausible that one emergent, specialised system in the brain could develop atypically while all those surrounding it develop normally. Atypicalities in one part of the system are likely to have ramifications on the developments of other parts of the system. While the isolated atypical development of individual functional components is not impossible, it would occur only under a narrow set of neurocomputational constraints. Thomas and Karmiloff-Smith (2002a) identified several candidate constraints of this type, including strong structure-function mappings, strong competition, early irreversible commitment, guided specialisation, and resource limitations. Unfortunately, those who postulate selective developmental deficits in adult phenotypes rarely argue for such constraints, let alone support them with empirical evidence. Instead, a maturational account is simply assumed that would support selective deficits in outcome.

A maturational account of uneven cognitive profiles in adult phenotypes would predict that the same profile should be found at earlier stages of development. Paterson et al. (1999) sought to test this hypothesis by comparing the disorders of Williams syndrome and Down's syndrome (DS). In the adult phenotype, WS demonstrates greater ability in language than DS, while DS demonstrates better ability in numerical cognition. This pattern was replicated by Paterson et al. using standardised receptive vocabulary tests and numeracy judgement tasks. However, when Paterson et al. explored the respective performance of toddlers with WS and DS using preferential looking measures to tap each domain, they found a different relative

profile. While both groups were very delayed, toddlers with WS and DS exhibited equal performance on a language task while toddlers with WS demonstrated superior performance to the DS group on a numeracy task. To the extent that the infant and adult tasks assessed the same aspects of the respective cognitive systems, then this study contradicts the notion that atypical cognitive profiles in infancy are miniature versions of those shown in adulthood. At the very least, the story involves differential delays and/or non-linear developmental profiles in the two disorders and requires a focus on development rather than just endstate deficits.

In addition, the idea that uneven cognitive profiles in genetic disorders can be explained by isolated, atypically developing functional brain systems does not fit well with what is currently known about how genes control brain development. Pennington (2001) summarises three broad classes of genetic control. These are effects: (1) on brain size, in terms of altering the number of neurons or synapses; (2) on neuronal migration, sometimes in a regionally specific fashion; and (3) on neurotransmission, either by changing levels of neurotransmitter or the binding properties of receptor proteins. Such genetic effects do not appear to operate in a region-specific fashion over the areas of cerebral cortex that eventually underlie higher cognitive processes (Kingsbury & Finlay, 2001). Regional specialisation is achieved by diffuse gradients of gene expression along with activity-dependent processes, although the primary sensory cortices and the limbic system are to some extent exceptions to this characterisation (see Kingsbury & Finlay, 2001, for discussion). That is to say, there are no current candidate genes that could influence, for example, the development of a language module without other albeit subtle widespread differences in brain development.

The case of the KE family is illustrative on this point. In this family, certain members demonstrated what was initially reported as a language-specific developmental deficit and the deficit was linked to the mutation of a single gene (FOXP2). However, subsequent detailed research on the family has revealed widespread structural and functional brain differences in affected family members, beyond those areas of the brain typically associated with language function in normal adults (e.g., Watkins et al., 2002). Moreover, other behaviour deficits, albeit of a subtler nature, have been found outside the domain of language, for example in performing less sophisticated oral-facial movements, and in nonverbal tasks involving rapid associative learning (e.g., Watkins, Dronkers, & Vargha-Khadem, 2002).

In line with the idea that developmental disorders do not involve region-specific structural atypicalities, post-mortem studies of genetic developmental disorders, and subsequently a growing body of work in structural brain imaging, have revealed widespread anomalies in gross and fine anatomy of the brains of these individuals. Gross anatomical differences can be found in disorders such as WS (Bellugi et al., 1999), DS (Nadel, 1999), and FraX (Reiss et al., 1995), in both the relative and absolute size of large-scale structures. Table 1 illustrates finer scale cytoarchitectonic and dendritic abnormalities found across a range of disorders (Kaufmann & Moser, 2000).

Table 1 (adapted from Kaufmann and Moser, 2000)

Neocortical cytoarchitectonic and dendritic abnormalities in genetic disorders associated with Mental Retardation [a]				
Disorder	Laminar disturbance	Increased packing density	Reduced dendritic length	Spine dysgenesis
Down syndrome	Y	N	Y	Y
Fragile-X syndrome	N	N	N	Y
Neurofibromatosis-1	Y (focal)	N	?	?
Patau syndrome	N	N	Y	Y
Tuberous sclerosis	Y (focal)	Y (focal)	Y (focal)	Y (focal)
Williams syndrome	Y	Y	?	?
Phenylketonuria	N	Y	Y	Y
Rett syndrome	N	Y	Y	Y
Rubinstein-Taybi syndrome	N	Y	?	?

[a] The conditions have been listed according to estimated incidence following Moser (1995)

Given the presence of widespread brain differences in many developmental disorders and given that, as we have argued, current evidence encourages the view that functional modules in the adult are not pre-specified in the infant but emerge as a product of development, it is clear that explaining uneven cognitive profiles in the adult phenotype of developmental disorders will be a complex endeavour. It appears likely that a final account of developmental deficits at the cognitive level will need to begin by identifying differences in low-level neurocomputational properties, perhaps in numbers of neurons and their thresholds, local or global connectivity, and activity-dependent changes in these parameters (Karmiloff-Smith, 1998; Oliver et al., 2000). The perturbations that these initial differences cause on the subsequent developmental trajectories of emerging cognitive systems must then be mapped out, taking into account atypical interactions, both internally between developing components and externally with the environment.

However, in terms of specificity of cause and outcome, our understanding of the relationship between neurocomputational parameters and cognitive performance is at present limited. For example, it might be possible that a computational property is

anomalous throughout the brain but only impacts on those cognitive domains that particularly rely upon it during development. Or it might be that the cytoarchitectonic properties that specify regions of cortex are disrupted by diffuse gene expression gradients in such a way that computational anomalies are topologically restricted (despite wider structural differences across the brain).[1] The latter possibility might support a more restricted scope for the cognitive domains impacted during development. Such issues remain to be worked out.

Nevertheless, FraX gives an indication of the possible complexity of the task ahead. FraX is associated with the silencing of a single gene, FMR1. Its gene product (FMRP) is normally involved in mechanisms of experience-dependent plasticity throughout the brain (Greenough and colleagues; e.g., Churchill et al., 2002, Greenough et al., 2001). However, alongside generalised delay, FraX exhibits an uneven cognitive profile in the adult phenotype. The complex interaction of FMRP with other proteins across development implicates a series of imbalances that have cascading effects on other elements of the developmental pathway at differing times through ontogeny (Scerif et al., in press; see Scerif, 2003, for a discussion). Thus, a brain-wide change at the cellular level may have specific outcomes via interactions across development. Specificity of outcome may be the result of temporally localised rather than just spatially localised events.

The potential for widespread effects in atypical functional brain development places certain requirements on the way in which behavioural deficits are studied in developmental disorders. Standardised cognitive tests that reveal behaviour in the

[1] For example, Karmiloff-Smith (1998) and Thomas & Karmiloff-Smith (2002b) discuss the case of a body-wide genetic difference which only impacts on hearing, resulting in the specific outcome of hereditary acquired deafness.

normal range cannot necessarily be assumed to demonstrate that the behaviour is being achieved by normal underlying cognitive processes, particularly if widespread deficits have more subtle effects on some cognitive domains than others. Sensitive and in-depth measures of apparently normal behavioural abilities in developmental disorders are therefore prompted by a neuroconstructivist approach (Karmiloff-Smith, 1998). In several cases such investigations have revealed normal-looking behaviours are achieved by atypical underlying processes (e.g., face recognition in WS: Karmiloff-Smith, 1997; Karmiloff-Smith et al., 2003; numerical processing in WS: Ansari et al., 2003). However, some of this work, particularly on face processing, remains controversial (see later).

If one compares developmental disorders with cases of early acquired brain damage in healthy children, it becomes readily apparent that the appropriate way to conceive of the disorders is in terms of the constraints that shape development rather than in the loss or impairment of specific cognitive structures (Karmiloff-Smith & Thomas, 2003a; Pennington, 1999). The most informative comparison is not a direct one, but via cases of adult acquired brain damage. The exercise works as follows. For behavioural deficits of adults with a given developmental disorder, identify which area(s) of the brain of a healthy adult would have to undergo focal damage for the individual to show this deficit. Then examine the consequences of early focal brain damage in otherwise healthy children, occurring to the same area(s) of the brain. What is the behavioural deficit exhibited by these individuals once they have reached adulthood? Does it match up with the deficits shown in the adult with the developmental disorder? In almost every case, the answer is the children with early brain damage show recovery and no lasting behavioural deficits (see Karmiloff-Smith

& Thomas, 2003a, for an example of this comparison carried out with regard to language in Williams syndrome). This begs the question, why do healthy children show recovery after early focal brain damage whilst individuals with developmental disorders that sometimes show apparently specific behavioural deficits do not? The answer is that the two cases constitute different limits on plasticity, i.e., differences in the ways that the healthy and atypical brain can be modified by experience (Thomas, 2003a).

Thomas (2003a) argues that a comparison of developmental disorders and children with acquired brain damage actually suggests that the closest similarities lie between developmental disorders and children with widespread early brain damage. In the latter group of children, recovery is limited and development increasingly diverges from the normal pathway with age (V. Anderson, Northam, Hendy & Wrennall, 2001). This comparison fits more closely with the widespread structural anomalies found in the brains of individuals with genetic developmental disorders.

Developmental disorders, then, strike at the heart of the issues we are considering in this book. What are the constraints that shape development? In the next section, we consider developmental deficits from the perspective of principles such as interactive specialisation, partial representations, timing, embodiment, and ensocialment. What happens when these constraints go awry?

Implication of constructivist principles

Interactive specialisation / embrainment

The interactive specialisation account of functional brain development argues that processing becomes both more localised and more specialised with development.

However, several developmental disorders suggest that this process may be deflected by atypical constraints.

Event-related potential (ERP) studies of face processing have indicated that upright and inverted faces elicit waveform components that differ both in amplitude and location on the scalp. When adults are presented with two matching faces vs. two faces that do not match, the ERP differences for upright faces in normal adults show a negativity around 320 ms that is largest over anterior regions of the right hemisphere. For inverted faces, however, the main difference between matched and mismatched stimuli is a symmetrical positive waveform component over parietal regions occurring between 400 and 1000 ms (Bellugi et al., 1999). When the equivalent waveforms for adults with WS were examined, three differences emerged (Mills et al., 2000): (i) the WS group exhibited the mismatch effect at 320 ms for both upright and inverted faces; (ii) the 320 ms waveform component did not show the right-hemisphere asymmetry of normal adults but was bilateral; (iii) there was an abnormally large negative wave component at 200 ms both to upright and inverted faces, which Bellugi et al. (1999) argued is linked to increased attention to faces in adults with Williams syndrome and appears specific to the disorder (see Grice et al., 2001, for similar results and a comparison to face recognition in autism). In short, in WS, ERP activity patterns in adulthood were consistent with neural processing of faces that was both less localised (bilateral instead of right lateralised) and less specialised (elicited by both upright and inverted faces instead of just upright faces, as well as by monkey faces and by other objects).

Above we argued that developmental disorders represent atypical limits on plasticity, such that development cannot compensate for early functional brain

damage in the way it appears able to in individuals with normal gene expression. However, this does not imply that there is no compensation in developmental disorders. According to the principle of embrainment, each brain system develops in the context of other brain systems. If an anomaly emerges across development in one system, it may well produce ramifications in other systems, perhaps ones that are recruited (atypically and perhaps less efficiently) to drive the behaviour of importance to the individual. Thus fMRI studies have demonstrated that adults with (phonological) developmental dyslexia demonstrate less activity in left posterior temporal-parietal areas compared to controls during listening and reading tasks that are phonologically demanding (Brunswick et al., 1999; Flowers et al., 1991; Paulesu et al., 1996; Rumsey et al., 1992; Shaywitz et al., 1998). But several of these studies also report increased activity in occipital and/or frontal regions that may reflect efforts to compensate for developmental impairments in phonological abilities (Casey, Thomas, & McCandliss, 2001), for instance with the use of additional visual strategies.

Timing

We have argued that cognitive change is crucially affected by the timing of events in brain development. Yet basic processes such as synaptogenesis and myelination have been found to show atypicalities in developmental disorders. Thus recent evidence from the PET imaging of neurotransmitter systems indicates that alterations in the plasticity of brain areas (as indexed by the numbers of particular types of synapses) may not follow the normal course in developmental disorders (Huttenlocher, 2002). Chugani et al. (1999) found a difference when comparing children with autism and

healthy controls. In the controls, serotonin synthesis capacity (which depends in part on the number of serotonergic synapses) in 5-year-old children was twice the adult value, subsequently decreasing back to the adult value following synaptic pruning. Children with autism, by contrast, had a lower serotonin synthesis capacity than controls at age 5, but the level steadily increased to 1.5 times the normal level by age 15, implying both delayed early synaptogenesis and then decreased synaptic pruning. Huttenlocher (2002) notes that this abnormal pattern has been found in the primary visual cortex of animals deprived of normally formed visual images during the system's early sensitive period, implicating activity-dependent processes in this abnormal marker of neuroplasticity.

In contrast to the preceding less-then-more pattern, Becker et al. (1986) found that dendritic arborizations in the visual cortex of children with Down's syndrome were paradoxically greater than normal early in infancy but then considerably less than normal by the age of 2 years. Becker et al. speculated that the initial overabundance might be a consequence of a compensatory response to the absence of adequate synapse formation. In many cases, DS is also characterised by a postnatal delay in myelination (Wisniewski, 1990). The delay is initially global but then manifests primarily in those nerve tracts that are myelinated late in development, such as the fibres linking the frontal and temporal lobes (Nadel, 1999).

It is hard to see how the effect of these abnormalities could be construed in terms of a genetic impairment to an isolated functional module. Alterations in synaptogenesis, aborization, and the myelination are inextricably linked to activity-dependent processes, that is, how the brain alters itself in response to experience.

Hierarchical integration / partial representations

Several developmental disorders have been characterised in terms of differences in the information that is abstracted from the environment prior to the operation of higher cognitive processes. Atypical partial representations of the environment can either facilitate or impair subsequent higher-level tasks. For example, in autism, individuals exhibit superior performance on visual search tasks compared to mental-age (MA) matched controls, where for instance a participant must pick out a green T in a field of red Ts (O'Riordan, 2000; O'Riordan, Plaisted, Driver & Baron-Cohen, 2001). O'Riordan and colleagues have argued that the superior performance arises not through attentional biases in higher processes but because individuals with autism begin by encoding greater discriminability between the components parts of visual scenes, thus facilitating selection in 'odd one out' tasks. In Specific Language Impairment, it has been argued that information about word sounds is represented in such a way that higher cognitive processes like inflectional morphology and syntax cannot operate as efficiently on word forms, particularly under time pressure (see Leonard, 1998, for discussion). Similarly, in developmental dyslexia, it has been argued that word sounds are represented in such a way that it becomes much harder to learn the association between the sounds of words and their written forms (See later section for detailed discussion of dyslexia). In general, alterations in the level of abstraction achieved in forming internal representations or in the dimensions of similarity that those representations encode can play a material role in the ability of other brain systems to employ these representations to drive additional processes. In the proposals on autism, SLI, and dyslexia, the consequence of atypical similarity

structure in lower representations is a deficit in processes higher up in a hierarchy of representational systems.

Embodiment

In many developmental disorders, the broad topology of the body is normal and the individual can perform common physical activities (although sometimes fine motor performance is reduced). In some disorders, however, movement can be restricted. This provides a potential window on the influence of the body on cognitive development. However, to date, there are few robust findings in this area. Although its findings are somewhat controversial and the conclusions speculative, one study serves to illustrate the directions such research may take, and the way in which atypical embodiment might impact on development.

Children with Spinal Muscular Atrophy (SMA) show physical weakness. Their first six-month's progression is normal, such that these children can sit unaided but they never achieve the ability to stand and walk. Sieratzki and Woll (1998) examined the language development of children with SMA. They exhibited normal vocabulary and use of irregular inflectional forms but over-regularisation, a marker for the acquisition of linguistic rules, was accelerated. Sieratzki and Woll speculated that the inability of these children to explore objects and forms in the environment might have advanced their analysis of patterning in language and the extraction of regularities. They speculated that at a neural level, the weakly used pre-frontal motor areas of the brains of children with SMA were being exploited by grammatical processing to accelerate developmental processes. Of course, an individual that has different physical abilities co-specifies a different effective environment. Thus, while

knowledge of many vocabulary items appeared to be developing normally, children with SMA nevertheless exhibited difficulties with certain vocabulary items, such as 'action' words, 'outside things', and 'places to go'.

Serious consideration of the body and the way it interacts with the world can lead us to re-conceptualise the problems that cognitive development must overcome. It may be that in the future, the study of atypical embodiment can clarify the extent to which development is constrained by a particular set of physical interactions with the world.

Ensocialment

The atypically developing child co-specifies an atypical environment. This interactive effect may be straightforward: a child with dyslexia may spend less time reading because it is a struggle, resulting in reduced input to the relevant cognitive systems. However, the interactions may be subtler, operating on the effective social environment to which the individual is exposed. Two studies exploring language development in DS and WS illustrate this idea.

The parents of children with developmental disorders involving learning disabilities are understandably concerned about the developmental progress of their offspring. Such anxiety may lead to changes in the effective social environment for the child. Thus Cardoso-Martins, Mervis, and Mervis (1985) found differences between the parental language input of children with DS compared to that of typically developing controls in terms of the language parents used to label objects for the child. While 67% of mothers of typically developing children used basic-level category terms to label objects in naming, only 31% of mothers of children with DS

used the basic-level category. Mothers of children with DS more frequently used precise object names (e.g., lion) than generic basic level terms (e.g., cat) during object labelling. It is possible that this was due to an increased concern that their children might not come to learn the correct names for objects spontaneously. There is no evidence either way on whether this strategy was beneficial, but it serves to show that atypical development cannot be considered solely from the perspective of the atypical brain but must extend to consider interactions with an atypical environment.

The effective social environment may be altered in more indirect ways by developmental deficits. For example, Laing and colleagues examined socio-interactive precursors to language development in toddlers with WS compared with MA matched controls (Laing, Butterworth, Ansari, Gsödl, Longhi, Panagiotaki, Paterson & Karmiloff-Smith, 2002). Toddlers with WS were proficient at dyadic (2-way) interactions with a caregiver and indeed sometimes exceeded the scores of MA controls due to persistent fixation on the caregiver's face (see also Bertrand, Mervis, Rice & Adamson, 1993; Jones et al., 2000). However, there was a marked deficiency in triadic interactions that incorporated an object. Specifically, toddlers with WS had difficulty switching attention from the caregiver's face to an object that was being referred to in communication via pointing, looking, and naming. Shared attention to newly named objects appears to be one of the main routes into vocabulary acquisition in normal development. The atypical nature of the social interaction found in children with WS may therefore have further ramifications for subsequent language development in this disorder. Indeed language development is delayed in this disorder. But in addition, there is accumulating evidence both that precursors to vocabulary development in WS are atypical (Karmiloff-Smith & Thomas, 2003a) and

that subsequent trajectory of language development is subtly different – despite the eventual relative strength of language in WS (Thomas & Karmiloff-Smith, 2003a).

Theories of both typical and atypical development will incorporate interactions with the social environment that the child co-specifies. A consideration of atypical development may additionally reveal the importance of these constraints in guiding normal development.

Relationship to theories of general cognitive variability

We have argued that one must view developmental disorders in terms of atypical neurocomputational constraints deflecting the normal path of development. However, this perspective raises wider issues about the nature of cognitive variability itself, of which atypical variability is only one kind. As we saw in Chapter X, there is a debate as to how learning and development are related, but each is undeniably a form of cognitive variation. Cognitive performance varies not just with age but also between individuals of the same age in the form of intelligence, and performance eventually declines in the form of aging. In addition, cognitive performance can show much variation on an hour-by-hour or minute-by-minute basis, in the form of arousal or in the form of attention. It is legitimate then to ask, what are the neurocomputational constraints that underlie each form of variation? Are the constraints that are altered in developmental disorders qualitatively different from those that are altered to explain other forms of variation? (See, e.g., M. Anderson, 1999 for one analysis of this question.)

This issue has been explored in more depth elsewhere, including a comparison of current proposals for the computational parameters that might underlie each type of

cognitive variation in developmental systems (Thomas & Karmiloff-Smith, 2003b). Intriguingly, however, recent work in behavioural genetics has indicated that cognitive variation in children at the extremes of the normal distribution (for instance, in the worst performing 5% on standardised tests) is more heritable than cognitive variation in the middle of the distribution (e.g., for measures of language development see Plomin & Dale, 2000). That is to say, extreme cognitive variation even within the normal sample may be under stronger genetic control and weaker environmental control that variation closer to the mean. There are interesting questions to pursue here. Consider three possibilities. (1) Perhaps for a given cognitive domain, a neurocomputational constraint acting within a certain range explains normal cognitive variation – but variation beyond this range leads to abnormal development. For example, say some notional parameter, "learning rate", takes values between 5 and 7 to account for normal variation. A rate of 4 or less would then produce a disorder. (2) Perhaps if the constraint shows extreme variation, it then reduces the potential of the environment to impose its own variability on performance. For example, if you have a learning rate as low as 3, it might then make less difference if the environment is rich or poor. The result of decreasing environmental influence and increasing influence of the initial parameter set would be a higher level of measured heritability for the disorder. (3) Or perhaps separate neurocomputational constraints exist for normal and atypical variability, with extreme normal performance and atypical development operating under the influence of different genes. For example, variations in "learning rate" would explain individual variation and be influenced by gene set X. A rather stupid system might have a low "learning rate" of 4. By contrast, in the normal case levels of another notional parameter, "computational resources" (influenced by gene

set Y) are fixed, say at level 10. If gene set Y is mutated in a disorder, "computational resources" may be reduced to 8, so producing cognitive impairments. Poor performance would be for different reasons than having a "learning rate" of 4. So far, these three possibilities have not been empirically distinguished. Theories on computational variability ideas are explored more fully in Thomas & Karmiloff-Smith (2003b).

In this context, we might also note that intelligence offers another example of the ensocialment of cognitive development, in the form of the Flynn effect. Flynn (1987, 1994) pointed out that mean IQ scores have been increasing each generation throughout the last century, with increases of between 5 and 25 IQ points measured across a large range of countries. Gains are more marked in real-time reasoning skills than tests of acquired knowledge. This generational gain seems paradoxical if intelligence corresponds to neurocomputational constraints of the developing brain that are in large part under genetic influence, supported by the continuing high heritability of IQ scores in twin studies. Indeed genetic relatedness can account for up to 75% of the variance in test performance (Jensen, 1973, 1998). However, Flynn has argued that the increase in test scores actually represents an increase in abstract problem solving rather than intelligence (or its neurocomputational correlates), and that the increase in this skill is not genetic but driven by an interaction with the environment. Dickens and Flynn (2001) have proposed a model whereby genetic differences that make individuals marginally better at abstract reasoning cause them to engage in additional practice and improvement of this skill. In their view, changes in patterns of work and leisure after World War II have placed a greater social emphasis on more cognitively demanding activities, triggering a "social multiplier"

effect whereby individuals compete to become ever more proficient at abstract reasoning. It is this competition that has resulted in the higher test scores. On the other hand, intelligence retains its apparently high heritability since it is the genetic inheritance (in terms of neurocomputational constraints that will affect cognition and the seeds of abstract reasoning ability) that pre-disposes the individual to expose himself or herself to a particular environment that subsequently exaggerates the cognitive variation.

Computational approaches to developmental disorders

Connectionist models of cognitive development form an ideal framework within which to explore the view that developmental deficits are the outcome of atypical neurocomputational constraints. Such models throw a particular spotlight on the role of initial computational constraints on the nature and success of subsequent trajectories of learning / development (Karmiloff-Smith & Thomas, 2003b). The ability of a model to acquire information from a given domain is limited by its initial architecture, activation dynamics, learning algorithm, and the representations through which the domain is depicted. In connectionist models of typical development, such design decisions are justified as far as possible via empirical evidence. A model is then judged successful if it captures the endstate competencies of the system as well as the developmental trajectory through which it passes. The opportunity here is to demonstrate that theoretically motivated alterations to the initial computational constraints of a normal model can then capture both the atypical trajectory and endstate behavioural deficits found in a particular developmental disorder. Where the success of a developmental model depends upon changes in the computational

constraints across development as in constructivist systems, then manipulations to the way in which such changes occur can also be explored as a candidate cause of developmental deficits (Westermann & Mareschal, 2003; Thomas & Karmiloff-Smith, 2002b).

The transition of a model from normal functioning to the disordered state is often the result of modifying quantitative variables, such as learning rate, levels of computational resources (processing units), or amount of noise. As such, connectionist models of developmental disorders lend themselves to an inherently continuous conception of pathology, with no absolute distinction between normality and disorder. However, alterations to models can be more radical, for instance using a different network topology or learning algorithm. These changes might be viewed as positing a qualitative distinction between normality and disorder. Referring back to our discussion on cognitive variability, the difference between these accounts will lie in the details of the developmental history of the processes that produced a computational system with these anomalies. Some parameters may alter quantitatively within the normal course of brain development, whilst others may require a genetic mutation to be altered.

Although this line of research is relatively new, connectionist models have already been used to explore the possible computational causes of deficits in several developmental disorders. Such investigations are contingent on the existence of valid models of typical development before parameter variations in the startstate (or rates of parameter change during development) can be explored. In consequence, work on atypical cognitive modelling tends to lag behind that on typical development. Developmental dyslexia, autism, SLI, and WS have all been the subject of recent

simulation work. Developmental dyslexia has been investigated by a number of researchers by manipulating the startstate parameters of models of reading development (e.g., Harm & Seidenberg, 1999; see later example for more detailed consideration of these models). Autism has been investigated by manipulating the startstate of models of category formation (Cohen, 1994, 1998). SLI and WS have been simulated by altering the startstate of models of inflectional morphology (Hoeffner & McClelland, 1993; Joanisse, 2000; Thomas & Karmiloff-Smith, 2003a). A review of the modelling of these disorders can be found in Thomas and Karmiloff-Smith (2002c) and in Mareschal and Thomas (2000). Evaluation of the contribution of connectionist modelling to the study of developmental disorders can be found in the peer reviews of two recent target articles, those of Oliver et al. (2000) and Thomas and Karmiloff-Smith (2002a). In the following paragraphs, we consider three examples of recent models with differing theoretical aims.

Oliver et al. (2000) examined the ways in which a process of feature map formation could be disrupted by changes in the initial properties of a self-organising connectionist network. They employed a neurobiologically constrained network in which a two-dimensional output layer received information from a single input retina. The network was presented with a set of stimuli in the form of bars lying across the input retina. Oliver et al. demonstrated that using their initial parameter set, the output layer formed a topographic map of the possible inputs: certain areas of the output layer specialised in responding to each input pattern and areas representing similar input patterns 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 for the output units, disrupted the connectivity

between the input and output layers, disrupted the connectivity responsible for lateral inhibition (competition) in the output layer, and changed the similarity of the input stimuli to each other. Importantly, these manipulations demonstrated that small differences in the initial constraints could have a very significant impact on the outcome of the developmental process. The resulting topographic map suffered a range of disruptions, including output units failing to specialise or simply turning off, specialisation emerging but not in organised areas, and organised areas emerging but without adjacent areas representing similar-looking bars. This model set out to illustrate a framework for considering the constraints that could perturb the normal trajectory of development. However, the model was not applied to any specific disorder, and importantly, the authors did not go on to demonstrate how, having developed a topographic map disrupted in a certain way, this led to deficits in a higher cognitive process which employed the map as input (Thomas, 2000; see Gustafsson, 1997, for a proposal of this nature with respect to autism).

In contrast, Thomas and Karmiloff-Smith (2003a) sought to capture the precise pattern of empirical data for a particular disorder and particular cognitive domain, but employed a connectionist model at a higher level of abstraction and therefore with fewer neurobiological constraints (Plunkett & Marchman, 1991, 1993). The disorder was Williams syndrome and the domain was English past tense acquisition. Individuals with WS have been reported to exhibit difficulties in generalising inflection patterns from words they know to novel items (Thomas et al., 2001) and in some studies with smaller subject numbers, selective difficulties with irregular past tenses (e.g., Clahsen & Almazan, 1998). Thomas and Karmiloff-Smith took a model of the acquisition of the past tense in normal development and altered

the initial constraints in line with empirical data concerning possible differences in phonological and semantic processing during language development in WS. The model was able to rule out certain candidate hypotheses for explaining the atypical developmental trajectory in the empirical data, such as solely delay (slower learning rate) and excessive internal resources (hidden units). In contrast, it established that phonological anomalies could lead to reduced generalisation of inflectional regularities to novel verbs, while semantic anomalies could produce difficulties in acquiring irregular past tense forms. In addition, the model demonstrated for the first time precisely how different computational constraints interact in a system during the process of development: the atypical trajectory found in WS past tense formation could arise from the combination of more than one altered constraint in the language system. Indeed, to explain the range of individual variation exhibited in the disorder, it appeared necessary to postulate multiple atypical constraints at work.

Thomas and Karmiloff-Smith (2002a) also used connectionist models to examine more general theoretical issues concerning the relation of developmental deficits to those found in cases of adult brain damage. This is another area to which connectionist models of cognition have been widely applied. Thomas and Karmiloff-Smith sought to assess whether disruptions occurring to the startstate of a learning system tended to produce the same deficits in performance as applying those same disruptions to the endstate of a normally trained model. The results of the modelling indicated that startstate damage to a system and endstate damage could in some circumstances cause similar behavioural impairments, but at other times the patterns were very different. The relationship depended on whether the system was able to use the developmental process to compensate for damage applied in the startstate, by

attenuating or even overcoming the effects of early anomalies. In other cases, early deficits followed by development produced worse deficits than damage to the end-state. Importantly, the simulations served to uncover the precise computational conditions under which each type of effect emerged. Moreover, the results convincingly demonstrated that in developmentally disordered systems, dissociations between impaired behaviour and behaviour-in-the-normal-range cannot be unambiguously interpreted without an understanding of the developmental conditions that pertained in the underlying system.

Thomas (2003b) recently pursued this issue further in a computational consideration of the multiple causality of behavioural deficits. Simulations indicated that narrowly defined behavioural deficits can potentially have multiple underlying computational causes. The implication is that developmental disorders defined on behavioural grounds alone (such as SLI or dyslexia) may gather together individuals with differing underlying cognitive architectures. This would seem to limit the ability of behavioural group studies to uncover any single 'cause' of the impairment defining the disorder. However, simulation work suggested that there may be behavioural markers for causal heterogeneity in the cross-measure variability within a disorder group. That is, the variability in performance can indicate the extent to which the atypical behaviour of a disorder group has a single or multiple underlying cognitive causes.

In short, computational models can help to explore the contribution of the developmental process to developmental deficits. They can assess the viability of claims concerning the possible origins of developmental deficits, and so begin to trace back these deficits to their genesis in early brain development. In the field of

developmental disorders, they serve to underline the crucial importance of formulating a precisely defined developmental account of a given cognitive ability before seeking to interpret behavioural deficits within a developmental disorder.

Challenges for the neuroconstructivist approach to developmental disorders

The neuroconstructivist approach to disorders offers the hope of building a bridge between genotype and phenotype via neurocomputationally constrained theories of cognitive development. However, a number of challenges lie ahead for this approach. In the following paragraphs, we consider four such challenges.

Challenge 1: When does normal-looking behaviour guarantee normal underlying process?

The first challenge stems from the claim that uneven cognitive profiles in the adult phenotype of developmental disorders should not automatically lead to a mechanistic explanation in terms of lists of "intact" (i.e., normally developed) functional modules vs. "impaired" (i.e., atypically developed) functional modules. The argument goes as follows. By the end of development, the atypical system will be different from a system that developed without impairment. Normal-looking behaviour may be produced by atypical underlying cognitive processes or atypical functional modules. The challenge here is to avoid assuming by definition that no normal-looking behaviour can be achieved by normal underlying cognitive processes in a developmental disorder. What are the conditions under which one will accept that a normal-looking behaviour is produced by a normal cognitive process?

In part, this is a methodological issue, revolving around the level of detail required in empirical investigation. For example, if an individual with a developmental disorder manages to produce a score 'in the normal range' on a cursory pencil-and-paper standardised test, one would be very hesitant in accepting this as evidence of normal underlying processes. The measure is unlikely to be able to discriminate behavioural outcomes produced by different cognitive processes. More sensitive tests are warranted, such as those involving measurement of real-time speeded performance and manipulation of implicit task variables (Tyler et al., 1997). However, there needs to be a point where score in the normal range in tests of a certain level of sensitivity are viewed as sufficient to establish normal underlying processing in a disorder. More colloquially speaking, eventually one has to be able to accept that if it walks like a duck and quacks like a duck, then it is most probably a duck. Specifying such a criterion, hard though it may be, is a valuable step: demonstration of the isolated, normal development of a given cognitive capacity in the face of a generally atypical system would be highly informative concerning the developmental constraints that must operate in acquiring the domain.

In part, however, this is also a theoretical issue. Claims of normality are sometimes made despite the fact that individuals with the disorder are not performing at the level that would be expected for their chronological age. A claim of normality can only be justified by arguing that 'delay' is in fact an explanatory concept. But there is very little mechanistic understanding of what causes developmental delay. Moreover, many disorders exhibiting 'delay' fail to reach normal adult levels however much time passes, arguing against simply slower development. Indeed such a

mechanistic claim would appear to make predictions about extended or absent sensitive periods during atypical development that have not been empirically pursued.

Challenge 2: To what extent do disorders exhibit atypical modularity?

The second challenge stems from the same claim, but now with regard to modular structure. If functional specialisation is a product of development rather than a precursor to it, it is possible that developmental disorders will exhibit cognitive systems with atypical functional structures. However, some researchers have argued that there is little empirical evidence to support this possibility. Thus Tager-Flusberg et al. (2003) maintain that there is much less deviance in the developmental processes and neurocognitive organisation in people with genetically based disorders than has been portrayed in the literature. Temple and Clahsen (2003) have argued that most behavioural data in developmental disorders appear amenable to characterisation in terms of normal modules that are under-developed, over-developed, or normally developed. Once more, there are both methodological and theoretical aspects to this challenge.

First, methodologically, few empirical tests are designed to assess whether underlying functional structure could be atypical. Standardised tests are often developed to test the integrity of normal cognitive components, and therefore data from such tests can only be interpreted in terms of those normal components (poor score = impaired component, normal score = normal component, high score = over-developed component, etc.). Sensitive tests that permit atypical responses are required to allow for the possibility of interpretations involving atypical underlying structures. Moreover, qualitatively atypical behaviours are apparent in some disorders. In autistic

savants who perform date calculation, this behaviour shows many of the hallmarks of an atypical functional module (fast, automatic, domain-specific, etc.). Individuals with synaesthesia demonstrate cross talk between sensory modalities so that, for example, certain noises can be reliably associated with certain colours. Such (objective) behaviours appear to implicate qualitatively atypical cognitive processes that allow, in this example, activation of visual cognition by auditory processes.

Alternatively, on theoretical grounds, it is possible that the range of behaviours that individuals can exhibit is constrained to some extent by the physical and social environment in which the individual's cognitive system is embedded. That is, behaviours normal or otherwise are in part constrained by the structure of the problem domain to which the cognitive system is exposed, whatever its underlying architecture. It is a serious issue of the extent to which cognitive architecture is visible in the behavioural changes and error patterns exhibited across development. The simplest example would be a cognitive domain that hard an easy part and a hard part. A wide range of learning systems would acquire the easy part before the hard part, and consequentially the developmental dissociation would tell us little about the actual learning system involved.

Computer simulations can give a more concrete demonstration of the influence of the task domain on developmental trajectories. In one set of simulations, the developmental profiles of qualitatively different processing architectures nevertheless exhibited strong similarities, stemming from the common problem domain to which all systems are exposed. Thus Figure 1 depicts the acquisition profiles for five networks with different architectures learning the English past tense. The architectures were a two-layer network, a three-layer network, a three-layer fully-

connected network, and two four layer networks with different amounts of internal units in each layer. Performance was assessed on the acquisition of regular and irregular verbs, levels of over-regularisation errors, and generalisation to different types of novel verbs (see Thomas & Karmiloff-Smith, 2003a). Whatever the architecture, certain clear developmental patterns were apparent in the respective trajectories, which emerged from the structure of the problem domain to which all networks were exposed. These were: (1) regular verbs were acquired more quickly than irregular verbs; (2) each type of network at some stage produced some degree of over-regularisation errors; and, (3) generalisation showed a similarity gradient depending on the relationship between the training and the test stimuli. Internal computational properties did modulate the developmental trajectories, but the modulation was bounded by constraints of the task domain. The message that this simulation exemplifies is that not all behavioural similarities between typical and atypical development must relate to similarities in cognitive architecture.

Of course, it could be that there is actually little variation in the modular structure of human cognition from the normal case to various types of developmental disorder. This would mean that the neurocomputational constraints guiding the emergence of functional modules were reasonably robust to atypical conditions. However, our understanding of the constraints that operate on the emergence of modular structure is at a relatively early stage, particularly with regard to atypical development (see Thomas & Karmiloff-Smith, 2002a, for discussion).

A couple of examples will suffice. On the one hand, Huffman et al. (1999) found that in marsupials, the entire projection field of the thalamus can come to be represented on a cortical volume that is substantially reduced in early development.

Huffman et al. removed the caudal one-third to three-fourths of the cortical sheet unilaterally at an early stage of development and examined the subsequent development and organization of the adult neocortex. Reduction of the size of the immature neocortex prior to the establishment of the thalamocortical connections was nevertheless associated with normal spatial relationships between visual, somatosensory, and auditory cortical fields on the remaining cortical sheet. Perhaps, then, given the role of activity in driving cortical specialisation, functional organization can follow a broadly normal pathway so long as thalamic input has not been disrupted by the genetic mutation in a given disorder, and normal conditions of competition exist in the remain portion of the sheet. On the other hand, as we saw earlier in electrophysiological studies of face recognition in individuals with WS, evidence already exists pointing to reduced localisation and specialisation of brain processes in developmental disorders. In sum, the actual atypicality of functional structures possible in developmental disorders remains to be clarified.

Figure 1. The effect of architectural changes on developmental profiles in a model of English past tense acquisition (Thomas & Karmiloff-Smith, 2003a). (h) = number of hidden units. Plots indicate the performance of the network across training on six metrics: (1) regular verbs (e.g., talk-talked); (2) irregular verbs (e.g., think-thought); (3) over-generalisation errors (e.g., think-thinked); (4) regular generalisation for novel verbs that do not rhyme with any existing irregulars (e.g., vask-vaske d); (5) regular generalisation for novel verbs rhyming with an existing irregular (e.g., frink-frinked, cf. drink); and (6) irregular generalisation for novel rhymes (e.g., frink-frank, cf. drink). [See Thomas et al. (2001) for equivalent empirical data in typical and atypical development.] Despite qualitative changes in architecture, developmental trajectories show strong similarities.

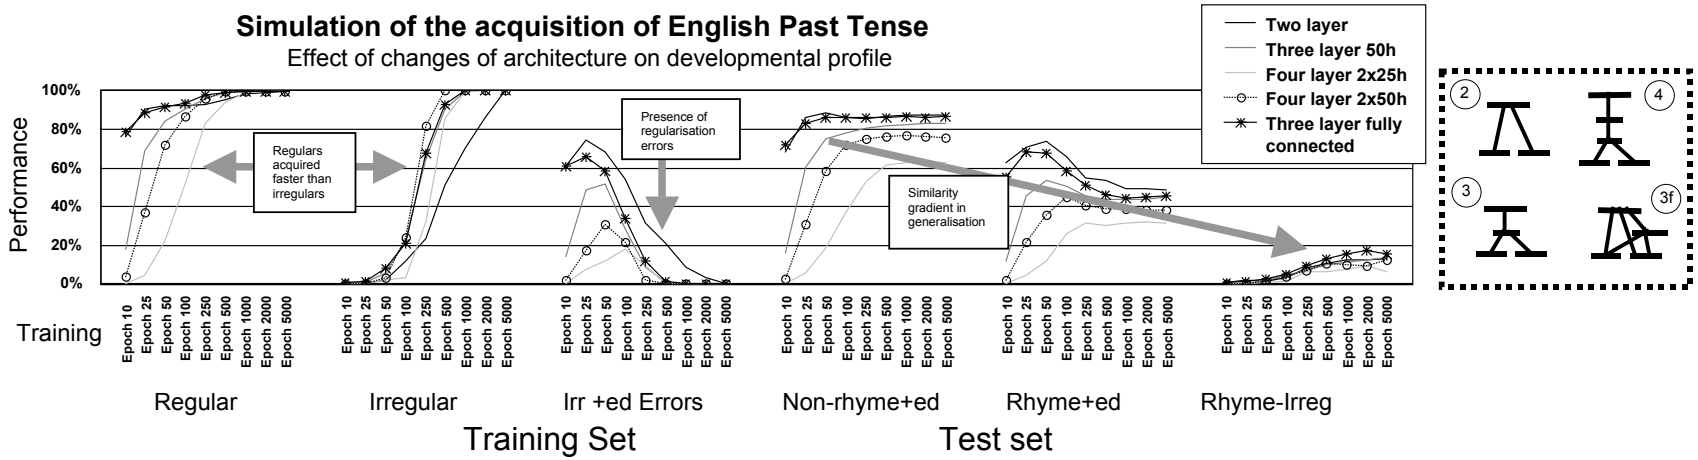

Challenge 3: How do variations in neurocomputational constraints map to variations in behaviour?

The third challenge for the neuroconstructivist approach is to work through the cognitive level implications of claims like the following: "subtle related initial deficits (e.g., firing thresholds which are either too high or too low) can give rise to huge differences in the end state which seem to bear no relation to one another" (Johnson, Karmiloff-Smith, Pennington, & Oliver, 2000, p. 38). It can be demonstrated that contrasting deficits in the endstate performance of developing systems can be produced by changes in a single initial computational parameter (e.g., see Thomas & Karmiloff-Smith, 2003a, for an example using the temperature of the sigmoid activation function). But what allows us to call the initial parameter difference "subtle" given that its ultimate impact is so significant? It cannot be that the numerical change in the parameter is itself small, or that we have changed only one parameter, for there is no absolute scale here.

There are two implications of this type of neuroconstructivist claim. Both are of theoretical importance but both ultimately requiring empirical support. First, "subtle" can mean that there is a non-linear relationship between changes in the startstate parameter and ultimate developmental outcomes. Thus, perhaps initial parameter changes across wide ranges produce little variability in the endstate, while much smaller changes in a sensitive range can produce great variability in the endstate. An emphasis on non-linearity is an important aspect of the neuroconstructivist approach (see Elman et al., 1996, chapter 4, for discussion). It is central in helping us understand how differences in the genotype might be related to

difference in the phenotype. Second, a "subtle" effect can mean that the developmental process itself exaggerates the impact of the parameter. In this case, if contrasting developmental profiles were to be found in the endstate of two disorders, one would expect much smaller behavioural differences in infancy. This points towards a particular empirical paradigm for comparing disorders across developmental trajectories, work that has begun to produce interesting results (see, for example, earlier discussion of work by Paterson et al., 1999).

Challenge 4: How do atypical brain structures relate to atypical cognitive structures?

The final challenge for the neurconstructivist approach is to understand what differences in the apparent structure or function of the brain in a given developmental disorder (as revealed, for instance, by brain imaging studies) in fact have computational consequences for the development of cognition. The difficulty here is that atypical functional structure at the cognitive level will produce atypical activation patterns in the brain, but atypical activations patterns do not seem to guarantee atypical functional structure. For example, 2-5% of 'normal' individuals appear to have right-lateralised language systems (Bates & Roe, 2001). Yet these individuals are not marked out as having atypical cognitive-level language systems. Women appear to demonstrate more bilateral patterns of brain activation in language tasks than men (e.g., Shaywitz et al., 1995; see Cameron, 2001, for discussion). Indeed, sex steroid hormones have been shown to modulate a wide range of brain processes including neurogenesis, cell migration, growth of the neuronal soma, dendritic growth, differentiation and synapse formation, synapse elimination, neuronal atrophy and apoptosis, neuropeptide expression, the expression of neurotransmitter receptors,

and neuronal excitability (Cameron, 2001). Yet, cognitive psychology does not (at present) posit qualitatively different functional structures for the language system in the two genders, let alone different overall cognitive architectures. Such differences in brain function are put down to the multiple realisability of cognitive architectures in neural structures, whereby the same cognitive level computations can be implemented in different ways in the wetware available.

Throughout this book, we emphasise the importance of incorporating brain-level constraints in theories of cognitive development. The negotiation between these two ideas – brain constraints altering cognitive architecture vs. multiple realisability – remains to be worked through. However, the outcome of this negotiation is likely to be highly influential in the evolution of future theories of cognitive development, and particularly in our understanding of the genesis of developmental disorders where brain constraints may vary.

References

Anderson, M. (1999). The development of intelligence. Hove: Psychology Press.

Anderson, V., Northam, E., Hendy, J., & Wrennall, J. (2001). Developmental neuropsychology: A clinical approach. Hove, Sussex: Psychology Press.

Ansari, D., Donlan, C., Thomas, M. S. C., Ewing, S., Peen, T., & Karmiloff-Smith, A. (2003). What makes counting count? Verbal and visuo-spatial contributions to typical and atypical number development. Journal of Experimental Child Psychology, 85, 50-62.

Baron-Cohen, S., H. Tager-Flusberg, and D. J. Cohen, Eds. (1993). Understanding Other Minds: Perspectives from Autism. Oxford: Oxford University Press.

Bates, E., & Roe, K. (2001). Language development in children with unilateral brain injury. In C. A. Nelson & M. Luciana (Eds.), Handbook of Developmental Cognitive Neuroscience (p. 281-307). Cambridge, Mass: MIT Press.

Becker, L. E., Armstrong, D. L., & Chan, F. (1986). Dendritic atrophy in children with Down's syndrome. Annals of Neurology, 20, 520-526.

Bellugi, U., Mills, D., Jernigan, T., Hickok, G., & Galaburda, A. (1999). Linking cognition, brain structure, and brain function in Williams syndrome. In H. Tager-Flusberg (Ed.), Neurodevelopmental disorders (pp. 111-136). Cambridge, Mass.: MIT Press.

Bertrand, J., Mervis, C., Rice, C. E., & Adamson, L. (1993). Development of joint attention by a toddler with Williams syndrome. Paper presented at the Gatlinberg Conference on Research and Theory in Mental Retardation and Developmental Disabilities, Gatlinberg.

Bishop, D. V. M., North, T., & Donlan, C. (1995). Genetic basis of specific language impairment. Developmental Medicine & Child Neurology, 37, 56-71.

Brunswick, N. McCrory, E., Price, C. J., Frith, C. D. & Frith, U. (1999). Explicit and implicit processing of words and pseudowords by adult developmental dyslexics: A search for Wernicke's Wortshatz? Brain, 122, 1901-1917.

Cameron, J. (2001). Effects of sex hormones on brain development. In C. A. Nelson & M. Luciana (Eds.), Handbook of Developmental Cognitive Neuroscience (pp. 59-78). Cambridge, Mass: MIT Press.

Cardoso-Martins, C., Mervis, C. B., & Mervis, C. A. (1985). Early vocabulary acquisition by children with Down syndrome. American Journal of Mental Deficiency, 90(2), 177-184.

Casey, B. J., Thomas, K. M., & McCandliss, B. (2001). Applications of magnetic resonance imaging to the study of development. In C. A. Nelson & M. Luciana (Eds.), Handbook of Developmental Cognitive Neuroscience (p. 137-147). Cambridge, Mass: MIT Press.

Chugani, D. C., Muzik, O., Behen, M., Rothermel, R., Janisse, J. J., Lee, J., et al. (1999). Developmental changes in serotonin synthesis capacity in autistic and non-autistic children. Annals of Neurology, 45, 287–295.

Churchill, J. D., Grossman, A. W., Irwin, S. A., Galvez, R., Klintsova, A. Y., Weiler, I. J., & Greenough, W. T. (2002). A Converging-Methods Approach to Fragile X Syndrome. Developmental Psychobiology, 40, 323-328.

Clahsen, H. & Almazan, M. (1998). Syntax and morphology in Williams Syndrome. Cognition, 68, 167-198.

Cohen, I. L. (1994). An artificial neural network analogue of learning in autism. Biological Psychiatry, 36, 5-20

Cohen, I. L. (1998). Neural network analysis of learning in autism. In D. Stein & J. Ludick (Eds.) Neural networks and psychopathology, pp. 274-315. Cambridge: Cambridge University Press.

Cornish, K.M., Munir, F., & Cross, G. (1999). Spatial cognition in males with Fragile-X syndrome: evidence for a neuropsychological phenotype. Cortex, 35, 263-271.

Cornish, K.M., Munir, F., & Cross, G. (2001). Differential impact of the FMR-1 full mutation on memory and attention functioning: a neuropsychological perspective. Journal of Cognitive Neuroscience, 13, 144-151.

Dickens, W. T. & Flynn, J. R. (2001). Heritability estimates versus large environmental effect: The IQ paradox resolved. Psychological Review, 108, 346-369.

Donnai, D., & Karmiloff-Smith, A. (2000) Williams Syndrome: from genotype through to the cognitive phenotype. American Journal of Medical Genetics, 97, 164-71.

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

Flowers, D. L., Wood, F. B., & Naylor, C. E. (1999). Regional cerebral blood flow correlates of language processes in reading disability. Archives of Neurology, 48, 637-643.

Flynn, J. R. (1987). Massive gains in 14 nations: What IQ tests really measure. Psychological Bulletin, 101, 171-191.

Flynn, J. R. (1994). IQ gains over time. In R. J. Sternberg (Ed.), Encyclopaedia of human intelligence (pp. 617-623). New York: Macmillan.

Frangiskakis, J. M., Ewart, A., Morris, C. A., Mervis, C. B., Bertrand, J., Robinson, B. F., Klein, B. P., Ensing, G., Everett, L. A., Green, E. D., Proschel, C., Gutowski, N. J., Noble, M., Atkinson, D. L., Odelberg, S. J. and Keating, M. T. (1996). LIM-kinase1 hemizygosity implicated in impaired visuospatial constructive cognition. Cell, 86, 59-69.

Freund, L., and Reiss, A. L. (1991). Cognitive profiles associated with the fragile X syndrome in males and females. American Journal of Medical Genetics, 38, 542-547.

Greenough, W.T., Klintsova, A. Y., Irwin, S. A., Galvez, R., Bates, K. E., & Weiler, I. J. (2001). Synaptic regulation of protein synthesis and the fragile X protein. Proceedings of the National Academy of Sciences, 98, 7101-7106.

Grice, S., Spratling, M.W., Karmiloff-Smith, A., Halit, H., Csibra, G., de Haan, M. & Johnson, M.H. (2001) Disorders visual processing and oscillatory brain activity in autism and Williams syndrome. NeuroReport, 12, 2697-2700.

Gustafsson, L. (1997). Inadequate cortical feature maps: A neural circuit theory of autism. Biological Psychiatry, 42, 1138-1147.

Harm, M. W. & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. Psychological Review, 106, 491-528.

Hoeffner, J. H. & McClelland, J. L. (1993). Can a perceptual processing deficit explain the impairment of inflectional morphology in developmental dysphasia? A

computational investigation. In E.V. Clark (Ed.), Proceedings of the 25th Child language research forum. Stanford: Stanford University Press.

Huffman, K. J., Molnar, Z., VanDellen, A., Kahn, D. M., Blakemore, C., & Krubitzer, L. (1999). Formation of cortical fields on a reduced cortical sheet. Journal of Neuroscience, 19, 9939-9952.

Huttenlocher, P. R. (2002). Neural plasticity: The effects of environment on the development of the cerebral cortex. Cambridge, MA: Harvard University Press.

Jackson, N. E., & Coltheart, M. (2001). Routes to reading success and failure. Hove: Psychology Press.

Jensen, A. R. (1973). Educability and group differences. New York: Harper and Row.

Jensen, A. R. (1998). The g factor: The science of mental ability. Westport, CT: Praeger.

Joanisse, M. F. (2000). Connectionist Phonology. Unpublished Ph.D. Dissertation, University of Southern California.

Jones, W, Bellugi, U., Lai, Z., Chiles, M., Reilly, J., Lincoln, A., & Adolphs, R. (2000). Hypersociability in Williams syndrome. Journal of Cognitive Neuroscience, 12: Supplement, 30-46.

Johnson, M. H., Karmiloff-Smith, A., Pennington, B. F., & Oliver, A. (2000). Deviations in the emergence of representations: themes and variations. Developmental Science, 3, 38-40.

Karmiloff-Smith, A. (1997). Crucial differences between developmental cognitive neuroscience and adult neuropsychology. Developmental Neuropsychology, 13, 513-524.

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

Karmiloff-Smith, A. & Thomas, M. S. C. (2003a). What can developmental disorders tell us about the neurocomputational constraints that shape development? The case of Williams syndrome. Development and Psychopathology, 15, 969-990.

Karmiloff-Smith, A. & Thomas, M. S. C. (2003b). Developmental disorders. In M. A. Arbib (Ed.) The Handbook of Brain Theory and Neural Networks, 2nd Edition. Cambridge, MA: MIT Press.

Karmiloff-Smith, A., Thomas, M. S. C., Annaz, D., Humphreys, K., Ewing, S., Grice, S., Brace, N., Van Duuren, M., Pike, G., & Campbell, R. (2003). Exploring the Williams Syndrome Face Processing Debate: The importance of building developmental trajectories. Manuscript submitted for publication.

Kaufmann, W. E., & Moser, H. W. (2000). Dendritic anomalies in disorders associated with mental retardation. Cerebral Cortex, 10, 981-991.

Kingsbury, M. A., & Finlay, B. L. (2001). The cortex in multidimensional space: Where do cortical areas come from? Developmental Science, 4, 125-157.

Laing, E., Butterworth, G., Ansari, D., Gsödl, M., Longhi, E., Panagiotaki, G., Paterson, S., & Karmiloff-Smith, A. (2002). Atypical development of language and social communication in toddlers with Williams syndrome. Developmental Science, 5, 233-246.

Leonard, L. B. (1998). Children with Specific Language Impairment. Cambridge, MA: MIT Press.

Mareschal, D. & Thomas, M. S. C. (2001). Self-organisation in normal and abnormal cognitive development. In A. F. Kalverboer & A. Gramsbergen (Eds.), Handbook

of Brain and Behaviour in Human Development (pp.743-766). Kluwer Academic Press.

Mills, D. L., Alvarez, T. D., St. George, M., Appelbaum, L. G., Bellugi, U., & Neville, H. (2000). Electrophysiological studies of face processing in Williams syndrome. Journal of Cognitive Neuroscience, 12:Supplement, 47-64.

Nadel, L. (1999). Down syndrome in cognitive neuroscience perspective. In H. Tager-Flusberg (Ed.), Neurodevelopmental disorders (pp. 197-222). Cambridge, Mass.: MIT Press.

Oliver, A., Johnson, M. H., Karmiloff-Smith, A., & Pennington, B. (2000). Deviations in the emergence of representations: A neuroconstructivist framework for analysing developmental disorders. Developmental Science, 3, 1-23.

O'Donnell, W. T., & Warren S. T. (2002). A decade of molecular studies of fragile X syndrome. Annual Review of Neuroscience, 25, 315-38

O'Riordan, M. (2000). Superior modulation of activation levels of stimulus representations does not underlie superior discrimination in autism. Cognition, 77, 81-96.

O'Riordan, M. A., Plaisted, K. C., Driver, J., & Baron-Cohen, S. (2001). Superior visual search in autism. Journal of Experimental Psychology: Human Perception and Performance, 27, 719-730.

Paterson, S. J., Brown, J. H., Gsödl, M. K., Johnson, M. H., & Karmiloff-Smith, A. (1999). Cognitive modularity and genetic disorders. Science, 286, 2355-2358.

Paulesu, E., Frith, U., Snowling, M., Gallagher, A., Morton, J., Frackowiak, R. S. J., & Frith, C. D. (1996). Is developmental dyslexia a disconnection syndrome? Evidence from PET scanning. Brain, 119, 143-157.

Pennington, B. F. (1999). Dyslexia as a neurodevelopmental disorder. In H. Tager-Flusberg (Ed.), Neurodevelopmental disorders (pp. 307-330). Cambridge, Mass.: MIT Press.

Pennington, B. F. (2001). Genetic methods. In C. A. Nelson & M. Luciana (Eds.), Handbook of Developmental Cognitive Neuroscience (pp. 149-158). Cambridge, Mass: MIT Press.

Pennington, B. F. & Smith, S. D. (1997). Genetic analysis of dyslexia and other complex behavioral phenotypes. Current Opinion in Pediatrics, 9, 636-641.

Plomin, R., & Dale, P. S. (2000). Genetics and early language development: A UK study of twins. In D. V. M. Bishop & L. B. Leonard (eds.), Speech and Language Impairment in Children: Causes, Characteristics, Intervention and Outcome (pp. 35-51). Psychology Press.

Plunkett, K. & Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. Cognition, 38, 1-60.

Plunkett, K. & Marchman, V. (1993). From rote learning to system building: acquiring verb morphology in children and connectionist nets. Cognition, 48, 21-69.

Reiss, A., L., Abrams, M. T., Greenlaw, R., Freund, L., & Denckla, M. B. (1995). Neurodevelopmental effects of the FMR-1 full mutation in human. Nature Medicine, 1, 159-167.

Rumsey, J. M., Andreason, P., Zametkin, A. J., Aquino, T., King, A. C., Hamburger, S. D., Pikus, A., Rapoport, J. L., & Cohen, R. M. (1992). Failure to activate the left

temporoparietal cortex in dyslexia, an H2150 PET study. Archives of Neurology, 49, 527-534.

Scerif, G. (2003). Infant and toddler precursors of attentional difficulties in fragile X syndrome: A neurodevelopmental perspective. Unpublished Ph.D. Dissertation, University of London.

Scerif, G., Cornish, K., Wilding, J., Driver, J., & Karmiloff-Smith, A. (in press). Visual search in typically developing toddlers and toddlers with fragile X and Williams syndrome. Developmental Science.

Sieratzki, J. S., & Woll, B. (1998). Toddling into language: precocious language development in motor-impaired children with spinal muscular atrophy. In A. Greenhill, M. Hughes, H. Littlefield & H. Walsh (Eds.), Proceedings of the 22nd Annual Boston University Conference on Language Development, Volume 2. (pp. 684-694). Somerville MA: Cascadilla Press.

Shaywitz, B., Shaywitz, S. E., Pugh, K. R., Constable, R. T., Skudlarksi, P. L., Fulbright, R. K., Bronen, R. A., Fletcher, J. M., Shankweiler, D. P., & Katz, L. (1995). Sex differences in the functional organization of the brain for language. Nature, 373, 607-609.

Shaywitz, S. E., Shaywitz, B. A., Pugh, K. R., Fulbright, R. K., Constable, R. T., Mencl, W. E., Shankweiler, D. P., Liberman, A. M., Skudlarski, P., Fletcher, J., Katz, L., Marchione, K. E., Lacadie, C., Gatenby, C., & Gore, J. (1998). Functional disruption in the organization of the brain for reading in dyslexia. Proceedings of the National Academy of Science (USA), 95, 2636-2641.

Simonoff, E., Bolton, P., and Rutter, M. (1998). Genetic perspectives on mental retardation. In J. A. Burack, R. M. Hodapp, & E. Zigler (Eds.), Handbook of

mental retardation and development (pp. 41-79). Cambridge University Press, Cambridge, UK.

Tassabehji, M., Metcalfe, K., Fergusson, W.D., Carette, M.J., Dore, J.K., Donnai, D., Read, A.P., Proschel, C., Gutowski, N.J., Mao, X. and Sheer, D. (1996) LIM-kinase deleted in Williams syndrome. Nature Genetics, 13, 272-3.

Tassabehji, M., Metcalfe, K., Karmiloff-Smith, A., Carette, M. J., Grant, J., Dennis, N., Reardon, W., Splitt, M., Read, A. P., & Donnai, D. (1999). Williams syndrome: Use of chromosomal micro-deletions as a tool to dissect cognitive and physical phenotypes. American Journal of Human Genetics, 64, 118-125.

Tager-Flusberg, H., Plesa-Skwerer, D., Faja, S., & Joseph, R., M. (2003). People with Williams syndrome process faces holistically. Cognition, 89, 11–24.

Temple, C. & Clahsen, H. (2003). How connectionist simulations fail to account for developmental disorders in children. Behavioral and Brain Sciences.

Thomas, M. S. C. (2000). Neuroconstructivism's promise. Developmental Science, 3(1), 35-37.

Thomas, M. S. C. (2003a). Limits on plasticity. Journal of Cognition and Development, 4(1), 95-121.

Thomas, M. S. C. (2003b). Multiple causality in developmental disorders: Methodological implications from computational modelling. Developmental Science, 6(5), 537-556.

Thomas, M. S. C., Grant, J., Gsödl, M., Laing, E., Barham, Z., Lakusta, L., Tyler, L. K., Grice, S., Paterson, S. & Karmiloff-Smith, A. (2001). Past tense formation in Williams syndrome. Language and Cognitive Processes, 16, 143-176.

Thomas, M. S. C. & Karmiloff-Smith, A. (2002a). Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. Behavioural and Brain Sciences, 25(6), 727-750.

Thomas, M. S. C. & Karmiloff-Smith, A. (2002b). Residual Normality: Friend or Foe? Behavioural and Brain Sciences, 25(6), 772-780.

Thomas, M. S. C. & Karmiloff-Smith, A. (2002c). Modelling typical and atypical cognitive development. In U. Goswami (Ed.), Handbook of Childhood Development (pp. 575-599). Oxford: Blackwells Publishers.

Thomas, M. S. C. & Karmiloff-Smith, A. (2003a). Modelling language acquisition in atypical phenotypes. Psychological Review, 110(4), 647-682.

Thomas, M. S. C. & Karmiloff-Smith, A. (2003b). Connectionist models of development, developmental disorders and individual differences. In R. J. Sternberg, J. Lautrey, & T. Lubart (Eds.), Models of Intelligence: International Perspectives, (p. 133-150). Washington, DC: American Psychological Association.

Tyler, L., Karmiloff-Smith, A., Voice, J. K., Stevens, T., Grant, J., Udwin, O., Davies, M., & Howlin, P. (1997). Do individuals with Williams syndrome have bizarre semantics? Evidence for lexical organization using an on-line task. Cortex, 33, 515-527.

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

Watkins, K. E., Dronkers, N. F., & Vargha-Khadem, F. (2002). Behavioural analysis of an inherited speech and language disorder: comparison with acquired aphasia. Brain, 125(3), 452-464.

Watkins, K. E., Vargha-Khadem, F., Ashburner, J., Passingham, R. E., Connelly, A., Friston, K. J., Frackowiak, R. S. J., Mishkin, M., & Gadian, D. G. (2002). MRI analysis of an inherited speech and language disorder: structural brain abnormalities. Brain, 125(3), 465-478.

Westermann, G., & Mareschal, D. (2003). Models of atypical development must also be models of normal development. Behavioral and Brain Sciences, 25, 771-772

Wisniewski, K. E. (1990). Down syndrome children often have brains with maturation delay, retardation of growth, and cortical dysgenesis. American Journal of Medical Genetic, Suppl 7, 274-281.