Chapter 12

Dyslexia: A case study of the application of the neuroconstructivist principles

So far, our case studies have focused on relatively low-level behaviours apparent in early infancy and that are generally culturally invariant. In this chapter we change tack substantially. We examine the emergence of reading and reading disorders. Reading is a cognitively complex task that involves low-level (perceptual) and high-level (cognitively hungry) operations. Focusing on reading allows us to illustrate the way in which the neuroconstructivist principles discussed in Chapter 5 interact at multiple levels well beyond infant development.

In this chapter, we consider the example of developmental dyslexia as an illustration of the principles of neuroconstructivism. In terms of *progressive specialization*, we will encounter the idea that the reading system emerges in part from specialization of the more general visual object recognition system that happens to have the correct properties for recognizing written words; environmental factors including instruction are essential in establishing this network. We will see evidence for reduced specialization in dyslexia. In terms of *competition and cooperation*, we will encounter the idea that the reading system divides labour between visual (word-specific) and phonological routes for naming; that the division of labour depends to some extent on the language in which reading is being learned; and that atypical neurocomputational constraints may disrupt the normal division of labour in the acquisition of reading, and in some cases provide opportunities for compensation. In terms of *partial representations*, we will see that different partial representations of a word are used to drive naming, and that developmental anomalies in different partial representations (phonological, orthographic) may lead to overlapping behavioural deficits in reading. In terms of *context-dependence*, we will see how new knowledge of orthography alters the structure of phonological representations at the onset of literacy. Additionally, in terms of *progressive specialization*, we will also see that phonological interventions for dyslexia are more effective

in aiding literacy development if they occur at the pre-school stage. Finally, we will see that the ease of learning to read depends on two factors: the granularity of the pre-existing phonological representations engendered by the individual's native language prior to onset of literacy, and the subsequent complexity of mapping from phonology to orthography for that language. The overt emergence of a disorder of literacy crucially depends on the latter complexity. Thus the *cultural specification* of the reading task influences both normal and atypical development. Research on typical and atypical reading development has produced a large and complex literature. In this chapter, we will once more consider convergent evidence from several disciplines, including psychology, neuroscience, genetics and computational modelling.

The reading system in normal development: the specific from the general

In discussing the cognitive and neural systems underpinning reading, it is important to remember that in historical terms, reading is a recent cultural invention reaching back just a few thousand years. Indeed, a significant proportion of humans still do not read. It is highly unlikely, then, that evolution will have had time to produce domain-specific constraints to guide the development of the relevant cognitive systems (see, e.g., McCandliss, Cohen, and Dehaene, 2003). If in an adult there is a reading system per se in the brain, then it must be a specialization of a more general system for relating visual forms to concepts and to speech. Such specialization will be driven by interaction with a particular kind of environment, namely, a culture that demands expertise in reading and mandates the investment of educational time in its acquisition.

The visual word form area

In the literate adult, brain-imaging research suggests that the reading system recruits up to a dozen different areas of the brain (Dehaene, 2003). Dehaene, Cohen and colleagues (e.g., Cohen *et al.*, 2000, 2002; Dehaene *et al.*, 2002; McCandliss, Cohen, Dehaene, 2003) have investigated one such region: the *visual word form area* in the left inferior temporal region and specifically the left fusiform gyrus. This area falls within the ventral stream typically involved in visual object recognition (and discussed in previous chapters). In imaging studies, the visual word form area has been found to be most active in response to written word forms. However, although activation levels are higher for orthographically legal letters strings than illegal letter strings, activation levels are equal for words and pseudowords (made-up letter strings that are plausible

words). Therefore the area appears to recognize visual stimuli that look like words prior to establishing whether they possess a meaning (Posner *et al.*, 1988; Dehaene *et al.*, 2002).

One has to be cautious in asserting a definitive functional specialization for a given area of cortex, since as Price and Devlin (2003) argue, the function of a given area may be defined in part by the other regions with which it is interacting in a given task. Nevertheless, in the silent reading of visually presented words, it has been argued that the location of the visual word form area is highly consistent across individuals, falling within approximately 5 mm irrespective of the language (and therefore script) that the individual has to learn (Dehaene, 2003; McCandliss *et al.*, 2003). McCandliss *et al.* (2003) argued that this consistency derives from the constrained nature of the task that script recognition demands of the visual system. They pointed to work by Malach and colleagues (2002) indicating the existence of a computational gradient across ventral visual cortex such that regions furthest from the center of the brain respond to fine-scale detailed images while those closer to the centre respond to larger scale images involving peripheral visual fields. The result is a gradient of areas that preferentially respond—moving respectively from the outer part of ventral visual cortex towards the centre—to objects and words, then faces, then buildings, and then outdoor scenes.

The area that eventually comes to specialize in the recognition of visual word forms is *the area that has the appropriate computational constraints for the recognition of visual stimuli* of this particular size and detail. These are the dual requirements for fine foveal discrimination and invariant recognition of letter- and word-sized units (McCandliss *et al.*, 2003). In addition, the area must comprise representations at a level of abstraction where the appropriate perceptual invariances required of script recognition can be computed. Fluent reading requires that the relevant partial representations of the visual input should be insensitive to changes in the font, size, case and position of words in the visual field.

Although the position of the visual word form area depends on the location in the brain where the appropriate information and computational constraints come together, Dehaene (2003) also argues that the range of constraints in the visual system of the pre-literate human brain would also limit the cultural variation of scripts that humans have invented. That is, human societies will only invent scripts to be written and read that are learnable given the constraints that the visual system contains. Humans and their cultural environment (including the current writing technology such as pens or keyboards) therefore co-specify each other via the external expression of the computational constraints of the cognitive system.

Nevertheless, the fact that a specialized reading function emerges across development from a more general system is underscored by brain imaging evidence that the activation of the visual word form area correlates with expertise in reading (Shaywitz *et al.*, 2002). In this study, activations levels were correlated with grapheme-phoneme decoding skills even when age was controlled for, implying that changes in the function of the visual word form area depend on expertise rather than genetic specification.

Cross-linguistic comparisons

While the computational properties of the human visual system have constrained the invention of written forms, humans have not always made things easy for themselves! Presumably for rather convoluted historical reasons, some languages selected by human cultures appear to be a lot harder to read than others. In alphabetic languages, the level of difficulty depends on two factors. The first is the appropriateness of the discriminations available in the *pre-literate* child's phonological representations, in terms of sound distinctions generated during acquisition of the language's spoken vocabulary, to which subsequent letters or letter clusters (graphemes) can map. The second is the consistency of the mapping between graphemes and component word sounds (phonemes) of a given language (Goswami, 2003). Let us consider these two factors in turn.

During language development prior to literacy, the requirement to learn increasing numbers of similar-sounding spoken word forms forces the phonological representations to acquire increasing levels of granularity (Metsala and Walley, 1998). According to the 'lexical restructuring theory', word forms are initially stored as undifferentiated wholes, but the acquisition of new similar (and therefore confusable) words pushes phonological representations to store words according to their component syllables (e.g., 'seesaw' => 'see-saw'). Subsequently, the acquisition of new similar (and therefore confusable) *syllables* pushes phonological representations to breakdown syllables into onset and rhyme components (e.g., 'see' => 's-ee', where *s* is the onset and *ee* the rhyme).

When children are required to learn to read, the task (particularly for reading novel words) is to learn the general relation between graphemes and phonemes. For languages like Italian and Spanish that have many simple consonant-vowel syllables, a decomposition of words into syllables, and syllables into onsets and rhymes is equivalent to splitting a word into component phonemes (e.g., 'casa' => 'c-a-s-a'). Children learning these languages are well placed to learn the relation of graphemes to components of their existing phonological representations. However, in other languages such as German

and English, many words have complex syllables ending in codas (e.g., the syllable 'cat' => 'c-a-*t*', where *t* is the coda). This means that with the onset of literacy, reorganization of phonological representations is required to a further level of granularity beyond onset and rhyme to distinguish codas as separate entities, before graphemes can be mapped to component phonemes (Goswami, 2003). However, not all languages use an alphabet—the Chinese writing system employs a large number of complex visual symbols representing units of meaning rather than units of sound. This represents a different challenge for the child learning to read, one that may rely more heavily on visual recognition skills (Vellutino *et al.*, 2004).

Turning to complexity, some alphabetic languages such as English and French have inconsistent mappings between phonemes and graphemes. Here *the context of the letters* surrounding a grapheme can alter the phonemes to which it corresponds. The amount of context may vary from the adjacent letters to the whole word. For example, in English, the letter 'i' in 'bit' maps to a different phoneme than the 'i' in 'bite', but to the same phoneme in 'sieve', but to a different phoneme in 'pint' and in 'aisle', and so on. This complexity presents additional challenges to deriving the relation between visual and spoken word forms, because the nascent reading system is required to develop and operate at multiple levels of granularity.

The consequence of the different computational challenges of each language is that cross-linguistic studies of reading have revealed differences in the rates at which literacy emerges. In all alphabetic languages, learning about letters appears to cause a reorganization of the mental lexicon into phoneme-based representations. For languages with consistent mappings between graphemes and phonemes, the reorganization is rapid, with grapheme-phoneme decoding ability close to ceiling within the first year of learning to read; in inconsistent languages like English, this process can take up to three years (Goswami, 2002, 2003). Thus, we see here a clear example of interactions between visual and auditory representations that are mediated by the complexity of the task present in the child's cultural environment.

Developmental disorders of the reading system: atypical constraints

Between 5 per cent and 17 per cent of the school age population exhibit a reading disability in English, depending on how the disorder is defined (Shaywitz and Shaywitz, 1994). Two main subtypes of developmental dyslexia have been identified (e.g., Castles and Coltheart, 1993; Manis *et al.*, 1996), although many dyslexics exhibit a 'mixed' pattern with symptoms of each

type. In *phonological* developmental dyslexia, there is particular difficulty in reading novel or pseudowords. In *surface* developmental dyslexia, there is a particular difficulty in reading irregular words (such as 'aisle', 'quay', and 'yacht' in English). These two patterns of reading deficit are also found in healthy adults after some kinds of brain damage (see Thomas and Karmiloff-Smith, 2002a, for discussion).

Genetics and dyslexia

There appears to be a significant genetic contribution to developmental dyslexia. This can be assessed by twin studies (see Pennington, 1999; Plomin and Dale, 2000; Plomin and Rutter, 1998, for discussion). While monozygotic (MZ) twins share the same genome, dizygotic (DZ) twins are on average no more genetically similar than siblings. To the extent that a developmental disorder has a genetic cause, one should expect MZ twins to be more likely to share the disorder than DZ twins. The 'concordance' rate of a disorder is defined as the proportion of twin pairs in which both members of the pair meet some cut-off criterion for possessing the disorder. When DeFries and Alarcon (1996) examined 200 pairs of MZ twins and 150 pairs of DZ twins in which at least one member of each pair met strict criteria for reading disability, they found MZ twins were 68 per cent concordant as compared to 38 per cent concordant for DZ. To the extent that twins tend to have much the same environment for learning to read, have roughly the same motivation to read and consequently, roughly the same level of experience with the task, the implication is that reading deficits are substantially due to genetic factors.

As discussed in Box 1.3 (p. 00), Heritability measures the proportion of the similarity between twins that can be attributed to their relatedness; the remaining similarity can be attributed to environmental influences either shared by the twins or unique to each twin (see, e.g., Gayan and Olson, 2001). Castles *et al.* (1999) investigated a large sample of twins with reading disability, classifying them according to whether they were more phonological or more surface dyslexics. Analyses suggested significant heritability for both subtypes, but stronger heritability for phonological dyslexics and a correspondingly higher contribution of the environment for surface dyslexics. This may be explained by the fact that surface dyslexia is measured by the ability to read exception or unusual words, which is itself dependent in part on exposure to printed matter. Subsequent analyses (Gayan and Olson, 2001, 2003) suggest a substantial overlap in the genetic influences on the ability to process words visually and phonologically in both typically developing and dyslexic readers.

AU: Please supply a page number for this cross-reference.

If the reading system is a recent, culturally induced specialization of more general systems involved in visual object recognition and spoken language processing, it follows that we should not expect a genetic variation that is entirely *specific* to the reading system. Nevertheless, stories in the media have appeared proclaiming the discovery of the 'gene behind dyslexia'.[1] One such story published in 1999 was prompted by work searching for quantitative trait loci (QTLs) correlating with developmental disorders (see Box 1.3, (p. 00 for a discussion of QTLs). QTLs are areas of chromosomal similarity in individuals who exhibit a developmental deficit (or individual variability on some trait) and represent candidate locations for genes implicated in contributing to the variability. In contrast to the idea that a disorder is caused by a single gene mutation not found in the normal population, the assumption behind QTLs is that many genes may contribute different quantitative amounts to the probability of having a given disorder, but that these genes may be involve in multiple functions (see Plomin and Dale, 2000, for an introduction). The media story was prompted by a report by Fagerheim *et al.* (1999) who studied a Norwegian family in which dyslexia was common and identified a region of chromosome 2 as a possible contributor. Other work has pointed to regions of chromosome 6 and chromosome 15 (see Cope *et al.*, 2005; Pennington, 1999; Schulte-Korne, 2001).

AU: Please supply a page number for this cross-reference.

The media story in question accepted that the processes involved in reading and writing might involve several genes. However, it viewed the identification of a gene for dyslexia as a potential breakthrough, since this would permit early screening and identification of dyslexics, followed by intervention to correct the impairment when the 'brains of children are most plastic', an idea to which we will return in the computational modelling section. Interestingly, the media story also quoted the chief executive of a national dyslexia association as arguing that the 'gene for dyslexia' shouldn't be 'removed or tampered with' because 'people with dyslexia, who are forced to think in alternative ways, are often revealed as geniuses'. The supporting empirical evidence for this claim comprised the case studies of a famous war leader, a famous businessman, a famous TV presenter, and a famous comedian, all of whom are dyslexics and all of whom are apparently geniuses!

A few years later in 2005, the same paper returned to the theme,[2] reporting that scientists had now discovered that a 'gene may be a cause of dyslexia'. Based on an association analysis (see Box 1.3, p. 00), Cope *et al.* (2005) found that variation in a gene on chromosome 6 (KIAA0319) was a reliable predictor of whether individuals had dyslexia or not. However, while it is known that the protein for which this gene codes is highly expressed in the brain, its function is as yet unknown. This reiterates the current gap

AU: Please supply a page number for this cross-reference.

that exists between genetic, neural, cognitive and behavioural approaches to disorders.

Finally, evidence for genetic effects should not be taken to underplay the influence of environment, particularly since reading requires both instruction and experience for its development. For example, Vellutino *et al.* (1996) carried out a longitudinal study on children classified in the first grade as reading impaired, a study that also included a remediation component. As many as two-thirds of the children were brought within the average range of reading ability after only one semester of remediation, while one third demonstrated persisting deficits particularly in phonological processing. This suggests that for two-thirds of the children, the initial reading deficit was experiential in origin. Where genetic effects do operate, gene–environment correlations are also likely to be the norm: parents with reading problems read less frequently to their children than parents without reading problems and have fewer books in their homes (Wadsworth *et al.*, 2000).

The specificity of cognitive differences and brain differences in dyslexia

In our discussion of developmental disorders in Chapter 11, we argued that genetic effects on brain development are rarely region-specific, and that in disorders of a genetic origin brain differences are frequently widespread. If the computational constraints of the developing brain have been altered by a genetic mutation in developmental dyslexia, we should not expect deficits to be reading-specific but to show up in other cognitive domains. Perhaps the reading task particularly exacerbates computational constraints that are anomalous in the wider visual system, affecting the ability to recognize visual word forms; or particularly exacerbates computational constraints that are anomalous in the spoken language system, affecting the ability to create the necessary granularity in phonological representations. Nevertheless, we should expect to find behavioural deficits elsewhere. The *specificity* of developmental dyslexia at both the cognitive level and the brain level, then, is highly relevant. However, currently there is little consensus in the literature at either level.

For example, at the *cognitive level*, it has been argued that individuals with dyslexia show deficits beyond reading and phonology, including various auditory deficits (Tallal, 1980), visual deficits (Lovegrove *et al.*, 1980), tactile deficits (Stoodley *et al.*, 2000), motor deficits (Fawcett *et al.*, 1996), learning deficits (Nicolson and Fawcett, 2000), and attention deficits (Hari and Renvall, 2001). Dyslexia has also been portrayed as a general sensorimotor syndrome (Stein and Walsh, 1997). However, some have argued that developmental dyslexia

is specifically phonological, with other features as co-incidental co-occurring factors or indirect side effects (e.g., Ramus, 2002; Ramus *et al.*, 2003).

As we shall see, causal models are complicated by the fact that the reading system is the emergent product of interactions between multiple systems. This means that a deficit in any one component may contribute to poor development in the whole. But if brain deficits are more widespread, some co-associated deficits may play no causal role but arise because adjacent areas of the brain have been affected during development. In their review of competing causal theories, Vellutino *et al.* (2004) argue that visual deficits (in the magnocellular system) may be just such a non-causal association of dyslexia. Longitudinal studies offer the best hope of tracing contributing factors, and have tended to support the idea that early difficulties with processing speech sounds predict later literacy problems (Pennington and Lefly, 2001; Snowling, Gallagher, and Frith, 2003). Domain-specific skills (phonological awareness, phonological memory visual-orthographic processing, and early literacy) are better predictors of later visual word recognition in the normal population than domain-general skills (general intelligence, higher-order reasoning, and language), but both domain-specific and domain-general skills predict similar amounts of variance in the more complex task of reading comprehension (Shatil and Share, 2003). However, the interactive nature of the reading system is confirmed by evidence that speech sound disorder (relating to intelligibility of production) and language impairment are *additive* risks for problems in pre-literacy skills in five–six year olds (Raitano *et al.*, 2004).

At the *brain level*, the last 5 years have seen around 50 scientific articles reporting the results of brain imaging experiments exploring brain anomalies associated with developmental dyslexia. The differences identified in the dyslexic brain show a significant degree of variability. In terms of functional magnetic resonance imaging (fMRI), the core dysfunction is claimed to reside in and around the angular gyrus of the left hemisphere, but with associated compensation in posterior regions in the right hemisphere and left inferior gyrus (Milne *et al.*, 2002). Claims are also made for anomalies in the organization of the cerebellum (Rae *et al.*, 2002), but so too for abnormalities of the magnocellular component of the visual stream specialized for processing fast temporal information (Stein and Walsh, 1997). On the basis of imaging evidence, dyslexics are claimed to differ from controls both in letter rhyming (phonological) and visual letter matching tasks (orthographic) (Temple *et al.*, 2001) and indeed not just in reading tasks but also tasks restricted to auditory language processing (Corina *et al.*, 2001). Simos *et al.* (2000) make the precise claim that dyslexia results from aberrant patterns of functional connectivity between ventral visual association cortex and temporo-parietal areas in left

hemisphere (see also Paulesu *et al.*, 1996). Structural imaging suggests that dyslexia is associated with marked rightward cerebral asymmetry and marked leftward asymmetry of the anterior lobe of the cerebellum (Leonard *et al.*, 2001), decreased grey matter in the temporal lobe (Eliez *et al.*, 2000), and lowered corpus callosum in posterior regions (Robichon *et al.*, 2000). McCandliss *et al.* (2003) pointed to evidence from several studies indicating a reduced tendency to activate the visual word form area during reading in dyslexics (fMRI: Paulesu *et al.*, 2001; Shaywitz *et al.*, 1998; MEG: Helenius *et al.*, 1999), arguing that this is evidence that there is an absence of *emergent specialization* of this region, through lack of relevant reading experience. In short, there is no consensus about whether structural differences are focused or widespread, or whether the range of tasks on which individuals show atypical brain activation patterns is narrow or diverse.

In terms of the two subtypes of developmental dyslexia, only phonological dyslexia has generated a consensus on the underlying cognitive cause. In the 'phonological representations' hypothesis (see Goswami, 2000; Snowling, 2000, for reviews), the initial representations of speech sounds that children develop are deficient, especially in terms of the emergence of the onset and rhyme level of granularity. When children come to learn to read, their phonological representations are ill fitted for learning the relationship between phonemes and graphemes, because the relevant speech sounds for the graphemes to connect to are not present, or at least their emergence is substantially delayed. As a consequence, the reading of novel words is impaired. Under this hypothesis, the problem pre-dates literacy, so it should also be apparent in tasks that do not involve visual (orthographic) information, such as short-term memory for words, non-word repetition, or speeded picture-naming tasks (e.g., Swan and Goswami, 1997); longitudinal studies of children at risk for dyslexia tend to support this view (Pennington and Lefly, 2001; Snowling *et al.*, 2003).

Importantly, the impact of initial phonological representations with reduced granularity has a differential impairment on the emergence of dyslexic symptoms depending on *the structure of the target language*. Thus, consistent languages like Spanish and Italian mitigate the symptoms, while inconsistent languages like English and French exaggerate the problems (Paulesu *et al.*, 2001). Nevertheless, non-orthographic tasks reveal spoken language deficits irrespective of language—for example, reading latencies are slower (Zeigler *et al.*, 2003). At the brain level, Paulesu *et al.* (2001) found that Italian individuals recruited on the basis of slower reading speed and defective phonological processing who nevertheless had high levels of reading accuracy demonstrated the same altered brain activation patterns as English and Italian individuals

with dyslexia. In this case, the common 'dyslexic' features were reduced activation in the left middle, inferior, and superior temporal cortex and in the middle occipital gyrus.

For surface developmental dyslexia, the picture is less clear. Studies have tended to agree that individuals with this subtype have failed to develop specific knowledge about word spelling, and that these individuals appear to demonstrate no deficits in their phonological representations, as assessed by tasks of phonological awareness (Castles and Coltheart, 1996; Goulandris and Snowling, 1991; Hanley *et al.*, 1992; Manis *et al.*, 1996; see Ans *et al.*, 1998, for discussion). Goulandris and Snowling (1991) found that individuals with surface dyslexia had poor performance on tasks of visual memory, although instances of surface dyslexia in the absence of visual impairments have also been reported (e.g., Castles and Coltheart, 1996). Since many of the reading errors diagnostic of developmental surface dyslexia have also been found in younger typically developing children (Bryant and Impey, 1986), and poor exception word reading can arise through lack of reading experience, some have doubted its validity as a distinct developmental dyslexia syndrome (see Pennington, 1999, for discussion). However, the view that surface dyslexia is simply delayed reading development fails to explain why delay should selectively strike the reading system. Moreover, as we saw in Chapter 11, the pattern of errors in a disorder is in part constrained by the nature of the task domain in any case, so that some similarities between early normal development and disordered development may be inevitable. The notion of 'delay' remains ill-specified at a mechanistic level and indeed, as we shall later see in Figure 12.2, computational modelling demonstrates that there are many constraints that produce slower learning as a side effect when those constraints become atypical.

Causal models of dyslexia

Let us assume for a moment that in older children with developmental dyslexia, we were to find anomalies that were not specific to brain areas involved in reading and behavioural deficits extending beyond the task of reading. At least two causal models are possible for how this state of affairs came about.

First, at the neural level, there could have been a reasonably restricted anomaly in brain development caused by genetic variation, followed by knock-on effects in the development of other brain areas. Alternatively, there could have been a more widespread effect of the genetic variation simultaneously affecting multiple brain areas. At the cognitive level, we could also have these two types of explanation: narrower-becomes-more-widespread-with-development

versus more-widespread-to-begin-with. At the cognitive level, however, there are the additional complications that some structural brain anomalies may have a greater or lesser *computational* consequence, and some cognitive processes may rely more or less on the computational properties that *are* affected by structural anomalies. The link between brain, neurocomputational properties, and behavioural deficits may be far from transparent.

In the face of these possibilities, the 'actual' cause of a reading deficit depends to some extent on whether there is only one way to disrupt the nascent reading system to produce phonological dyslexia and only one way to produce surface dyslexia. If there are multiple routes to each (as indeed will be suggested by the later computational modelling section), then subtypes may be causally heterogeneous. That is, experimental groups may be clusters of individuals drawn together by virtue of sharing a particular (impaired) *behaviour* (Thomas, 2003b).

A recent proposal allows us to contrast the two causal models. Ramus (2002, 2003) suggested that at a cognitive level, phonological developmental dyslexia is a consequence of a specific phonological deficit. However, in addition the deficit is sometimes accompanied by a sensorimotor syndrome with variable manifestations but little additional impact on phonology and reading. Under this hypothesis, the original cause at the neural level are focal anomalies of neural migration in the outer layer of the cortex, which are located mainly in the left peri-sylvian areas (Galaburda *et al.*, 1985; Humphreys *et al.*, 1990). These focal migration anomalies are viewed as being of genetic origin, based on evidence from mouse studies (Sherman *et al.*, 1990). However, *secondary* to the cortical anomalies, thalamic anomalies arise in the magno-cells of the lateral and medial geniculate nuclei of the thalamus, which produce visual and auditory deficits respectively (Livingstone *et al.*, 1991; Galaburda *et al.*, 1994). Based on animal studies, Ramus argued that this secondary effect only arises when the cortical anomalies coincide with excessive testosterone concentrations during brain development (Fitch *et al.*, 1997; Herman *et al.*, 1997; Rosen *et al.*, 1999). This would explain why behavioural genetic studies indicate that the phonological deficit in dyslexia is highly heritable, while auditory and visual deficits are not (e.g., Bishop *et al.*, 1999). Ramus invoked this account to explain why in a group of 16 adults with developmental dyslexia, a careful battery of tests revealed that all 16 had a phonological deficit, but 10 also had an auditory deficit, 4 a motor dysfunction, and 2 visual problems. Five, however, had a phonological deficit alone (Ramus *et al.*, 2003).

This account, linking a behavioural deficit with early genetic effects on brain development is speculative. Nevertheless, it illustrates our first causal model: that the initial genetic deficit is restricted to left peri-sylvian areas underlying

subsequent phonological processing, but under some conditions there are secondary effects on brain development that spread the impact of the initial anomaly. Ramus' theory currently lacks two key components: an account of the consequences of the ectopias and microgyri for the *computational* properties of peri-sylvian areas, and a developmental account of why an attempt to acquire phonological representations in a system with these anomalous computational properties should lead to the deficits we see in dyslexia. At it stands, there is a large leap from focal neural migration anomalies in the neonatal brain to adults who show reading impairments.

Ramus' (2003) hypothesis also highlights two of the issues raised in our earlier general discussion of developmental disorders. First, we argued that developmental disorders need to be viewed in terms of atypical limits on plasticity, given that healthy children with focal lesions do not usually show domain-specific deficits when they are older but instead exhibit recovery to the normal or low-normal range (Thomas, 2003a). If early brain anomalies in a dyslexic are indeed as focal as Ramus' account maintains, an explanation of subsequent behavioural deficits needs to address why there is no compensation of the cognitive system from other undamaged brain areas across development, sufficient to allow recovery.

Thomas and Karmiloff-Smith (2002a) identified several candidate computational constraints that would explain the isolated atypical development of individual functional components against a background of normal function, a condition they labelled 'Residual Normality'. Identification of such constraints is crucial if apparently selective behavioural deficits and an absence of compensatory recovery are to be fitted into a developmental theory of disorders. The computational constraints that would generate Residual Normality are as follows:

1. Strong structure–function correspondences—i.e., only one component has the properties to do the job, but it's malformed,
2. Strong competition,
3. Early irreversible commitment,
4. Inflexibly guided specialization—i.e., in 2–4: only one component is provided with the information to do the job, but it's malformed, and
5. Resource limitations—i.e., other components could do the job but they're busy.

In order to explain the mapping between early focal deficit and domain-specific end state deficit, Ramus (2002, 2003) appealed to two of the RN-preserving constraints, strong structure–function correspondences and strong

competition. He proposed first that certain anatomical modules are unique in their ability to support certain cognitive modules—i.e., only peri-sylvian areas have the computational structural properties to compute the functions required in the domain of phonological processing; and secondl that these areas are unique in their initial input and output connectivity—no other areas can compete to take over the input–output mappings of phonological processing (even if they had broadly appropriate computational properties) because initial connectivity is strongly biased against them. They are not provided with the relevant inputs and outputs to start to compete to perform the function. Together, these two constraints would ensure that no other area could effectively compensate for inefficiencies in peri-sylvian areas.

Of course, such a claim does not imply that individuals with dyslexia must demonstrate zero compensation compared to the normal reader in *attempting* to overcome phonological processing difficulties. Thus, we saw in Chapter 11 how brain-imaging studies of dyslexics revealed differences compared to controls in the activation of both occipital and frontal areas during reading tasks, consistent with attempts of other brain areas to compensate for inefficient phonological processing (see Casey *et al.*, 2001). But by virtue of being *defined* as dyslexics, we know that any such compensation must be insufficient to produce reading behaviour in the normal range. If individuals who had initial focal neural migration anomalies in peri-sylvian areas (or some other brain deficit) achieved successful compensation, then they would not have appeared in the dyslexic samples in any of these studies.

Longitudinal studies of children at risk for dyslexia have generated important insights (e.g., Pennington and Lefly, 2001; Snowling *et al.*, 2003). These children are identified because they either have a parent or an older sibling who exhibits dyslexia. The children can then be followed to see if they develop the disorder. Snowling *et al.* (2003) found that 66 per cent of their at risk children identified at 3 years and 9 months had developed dyslexia by 8 years of age, compared to only 13 per cent of a control group. However, the at-risk children who did not later develop dyslexia nevertheless showed poorer phonological processing skills than the control group at age six. This suggests that dyslexia is a continuous rather than an all-or-nothing disorder. Second, this at-risk non-dyslexic group showed higher verbal and performance IQ than the at-risk dyslexic group: compensation in reading development is therefore possible using other skills and cognitive components. Snowling *et al.* (2003) argued that one of these skills might be strong vocabulary, and therefore that problems in establishing a phonological pathway can be compensated for in an adaptive, interactive system by variability in other components.

This brings us to the second issue raised in our earlier general discussion of developmental disorders. In Chapter 11, we also introduced the thorny question of the relationship between atypical variability in developmental disorders and other forms of cognitive variability. In this case, the key issue is the relationship between individual variability (general and specific 'intelligences') and atypical development. In developmental disorders that are defined on behavioural grounds alone, the sample of individuals recruited into a study necessarily conflates individual variation with atypical variation due to the disorder. For example, let us assume that individuals show independent variation in their levels of auditory, visual, and motor abilities, simply as part of the individuation variation present in the normal population. Assume, too, that developmental dyslexia were to be caused (*contra* Ramus) by a genetic mutation with widespread effects in early brain development across auditory, phonological, visual and motor areas. However, in this scenario, certain dyslexic individuals with above average performance in one or other of these abilities could show behaviour in the normal range on some subset of visual, motor and auditory processing. (They could not show performance in the normal range in phonological processing because, by definition, they are recruited as dyslexics). The consequence would be individuals unified by sharing a phonological deficit, but varying in the other deficits they exhibit. This hypothesis would explain the same pattern of empirical data reported by Ramus *et al.* (2003), but appeal to the alternative causal model we identified earlier, that is, the simultaneous presence of more widespread genetic effects on brain development.

In terms of individual differences, low general intelligence can produce poor reading scores, but in this case against a background of low performance in other cognitive domains. It remains to be seen whether a reading deficit that is part of a domain-general pattern should be explained in terms of variations in different underlying neurocomputational parameters than in the case of domain-specific reading disability. Do individual variation and atypical variation converge on the same behavioural deficit? The same cognitive deficit? Fletcher *et al.* (1999) have argued that poor reading through generally low IQ and poor reading as an apparently domain-specific disability do not differ radically in the cognitive factors with which they are associated, such as poor phonological processing. A full account of cognitive variability in the developmental realm must one day unify domain-general and domain-specific accounts of poor performance (see Thomas and Karmiloff-Smith, 2003b, for discussion of this question).

Although Ramus' (2002, 2003) hypothesis omits a story of how initial computational deficits can lead to behavioural problems across development, other

researchers have been active in applying computational models to addressing this question. We now turn to consideration of these models.

Computational modelling of developmental dyslexia

The importance of constructing *developmental* models of reading is highlighted by insufficiencies in static, hand-wired computational models of the adult reading system. For example, one such model proposes two mechanisms for reading, one that relies on information about whole words, linking whole visual word forms to their pronunciations (the lexical route) and another that relies on finer granularity information, linking graphemes to their corresponding phonemes (the non-lexical route) (Coltheart *et al.*, 2001). If you attempt to extend this static model to explain developmental dyslexia, one story you can end up with is as follows:

> Some children might be acquiring the components of the lexical route at a normal rate, but be having difficulty with [developing] one or more components of the nonlexical route. Such children would have a selective difficulty in reading nonwords aloud. This is developmental phonological dyslexia . . . Other children might be acquiring the components of the nonlexical route at a normal rate, but be having difficulty with [developing] one or more components of the lexical route. This is developmental surface dyslexia.
>
> (Coltheart *et al.*, 2001: 246)

As we have seen, this sort of proposal begs the question, Why is there no compensation in the disorder? If only one of the two mechanisms is initially compromised, why doesn't the other initially intact mechanism compensate for the first mechanism across development? Indeed, Thomas and Karmiloff-Smith (2002a) demonstrated that dual-mechanism computational systems can exhibit precisely this sort of compensation unless they are prevented from doing so. Coltheart *et al.*'s proposal would only work if the reading system demonstrated residual normality (see previous section), a property that depends on the precise specification of the developmental process. Because it is a static model without a developmental process, the question is unanswerable.

By contrast, a large body of computational modelling work has attempted to explore the causes of reading deficits by implementing atypical developmental processes. We discuss these models in the following paragraph. First, however, a couple of caveats. Thus far, few of the models of developmental dyslexia have been based on a model of reading acquisition that captures all stages through which children pass when learning to read. Currently the normal models of development are somewhat limited. They tend to omit the early stages of the reading process when the child is using partial visual cues to recognize

whole words, and when the child's visual system has yet to reach the correct perceptual invariances required to recognize letters. For example, p, q, b, and d are confused because the child has to learn that normal perceptual invariances for visual objects such as rotation should not apply in the specialist domain of recognizing letters (Dehaene, 2003). Nevertheless, the existing models are sufficient to give an idea of the implications of various computational constraints in learning the relation between written and spoken words. The second caveat is that to date, there has been a relative lack of cross-linguistic modeling that would permit an exploration of the interaction of atypical computational constraints with the particular language to which the system is exposed. The final caveat is that, due to lack of space, this discussion omits a consideration of the relative merits of the competing *normal* models on which atypical manipulations are based. The reader is directed to discussions in Ans *et al.* (1998), Coltheart *et al.* (2001), and Harm and Seidenberg (2004).

Connectionist models of the typical and atypical reading system

Developmental models of reading have tended to appeal to connectionist architectures, employing two-layer and three-layer networks, usually with distributed representations. Some models have included attractor networks and cycling activation, so that the system settles into stable solution states.

Connectionist models of reading assume that the computational problem in this domain is to learn to map between representational codes of the written form of a word, the spoken form of a word, and the word's meaning (Plaut *et al.*, 1996; Seidenberg and McClelland, 1989). Typically, this involves three connectionist networks, one to map from orthography to phonology, one to map from orthography to semantics, and one to map from semantics to phonology, although in many models only the first of these networks is implemented (see Harm and Seidenberg, 2001, 2004, for an exception). Usually, each of these networks is assigned a three-layered structure comprising an input layer, an output layer, and an intermediate layer of hidden units.

Some models employ recurrent connections to allow phonological representations to settle into stable output (attractor) states, while a fully implemented network might allow cycling interactions between all three representational codes. Sometimes a layer of 'clean-up' units is connected to the phonological layer to aid the settling process (see e.g., Harm and Seidenberg, 1999, 2001). In other cases, orthography is connected to phonology via two routes, one with direct connections the other involving an intermediate hidden layer (Zorzi, Houghton and Butterworth, 1998a). In such a dual route network, the direct connections favour computations based on components of word

representations (graphemes and phonemes) while the indirect, hidden-unit-mediated route favours computations based on a larger granularity (combinations of graphemes or combinations of phonemes) (see Thomas and Karmiloff-Smith, 2002a, for discussion of this architecture and emergent specialization). Ans *et al.* (1998) used a similar two-route architecture but with localist hidden unit representations (in their terminology, a multiple-trace episodic memory). A localist representational format is one where the activity of a single unit corresponds to the representation of a single entity and there is no similarity between the representations of each entity. In the Ans *et al.*, model, the localist units were separately constrained to represent word components or whole words. Mappings between orthography and phonology could independently use one or other of these sets of processing units, prespecified as two processing routes (or 'modes').

A composite architecture summarizing all the various connectionist models is depicted in Figure 12.1. In line with the general theoretical framework for the

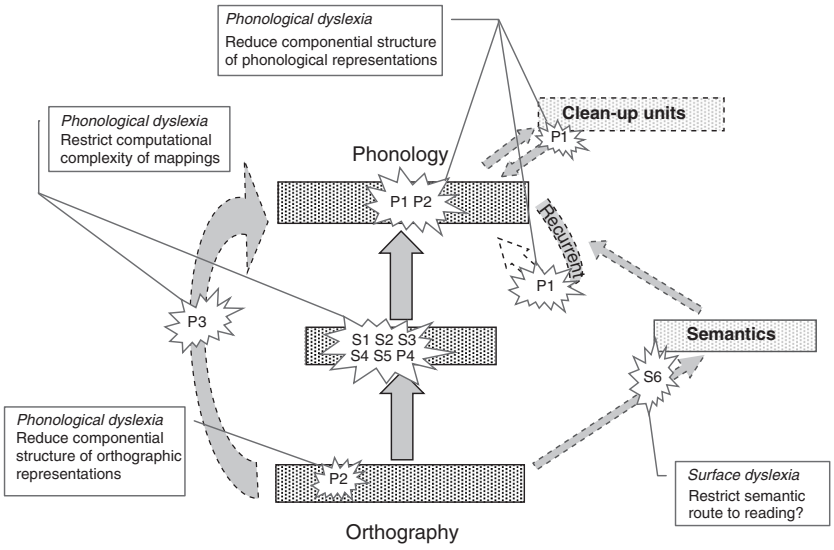

Fig. 12.1 Composite architecture of connectionist models of the reading system. The core assumption is a network mapping between orthography and phonology. Different models add other components (dotted outlines), such as (1) a semantically mediated route between these codes; (2) direct vs. indirect connections from orthography to phonology; (3) attractor networks for each representational code, include recurrent connections at output and/or clean-up units. P# stands for the site of start state manipulations designed to simulate developmental phonological dyslexia. S# stands for the site of start state manipulations designed to simulate developmental surface dyslexia (see text for references).

AU: Please supply a page number for this cross-reference.

simulation of developmental deficits outlined in Box 11.2 (Chapter 11, p. 000), the sub-types of dyslexia were simulated in these models by altering certain initial computational constraints and then exposing the system to a training set of written words and their pronunciations. Figure 12.1 also demonstrates the constraints that were changed in various models.

Simulating the subtypes of dyslexia

Surface dyslexia, an impairment in reading exception words, has been simulated by altering any initial constraints that reduce the general ability of the network to learn the relation between orthography and phonology (manipulations are labelled as 'S#' in Figure 12.1). Exception words are the first to suffer from this degradation, since they are inconsistent with most of the knowledge gained from exposure to reading words. Constraints that have this effect have included a reduction in the initial number of hidden units in this network (S1: Bullinaria, 1997; Harm and Seidenberg, 1999; Plaut *et al.*, 1996; Seidenberg and McClelland, 1989), a less efficient learning algorithm (S2: Bullinaria, 1997), less training (S3: Harm and Seidenberg, 1999), and a slower learning rate (S4: Harm and Seidenberg, 1999). In models that employ two routes to connect orthography and phonology, proposed manipulations have focused on impairing the indirect (hidden-unit-mediated) route (S5: Ans *et al.*, 1998; Zorzi, Houghton and Butterworth, 1998a). This route is able to process the larger granularity mappings required to encode whole word exceptions. Eliminating it reduces the network's ability to learn such exceptions using the remaining direct orthography–phonology route.

Plaut *et al.* (1996) argued that reading may take place either via connections between orthography and phonology, or via a semantic route. They argued that a *division of labour* would be negotiated between these two routes by a competition during learning (unimplemented in their model). They speculated that exception words (especially of a low frequency) might preferentially rely on the semantically mediated route. Snowling, Gallagher and Frith (2003) later proposed that some children at familial risk for dyslexia might be able to compensate for problems in the phonological pathway by making greater use of the semantic pathway, so long as their oral language skills were strong enough.

Simulations in the related domain of inflectional morphology, which is also characterized by a partially regular mapping problem, support the idea that exception mappings may preferentially rely on word-specific mechanisms. Joanisse and Seidenberg (1999) and Thomas and Karmiloff-Smith (2003a) have demonstrated that where word-specific information such as a meaning is

available at input, networks learning partially regular input–output mappings between uninflected and inflected phonological forms will exploit the word-specific information during learning to *support exception mappings but not regular mappings*. Moreover, Thomas and Karmiloff-Smith (2003a) found that removal of word-specific information from the startstate produced differentially delayed acquisition of the exception mappings compared both to regular mappings and generalization. It therefore seems likely that in the general reading architecture in Figure 12.1, a processing impairment (in connectivity, speed of processing, or quality of representations) in the semantically mediated route would also produce symptoms of developmental surface dyslexia compared to a normal system able to exploit this route (labelled S6 in Figure 12.1). Without the semantic route, the additional resource to support exceptions is no longer available. Indeed, simulations in the past tense domain suggest that an absence of word-specific information impairs low frequency exceptions more than high frequency exceptions, thereby exaggerating the frequency by regularity interaction—a pattern sometimes reported in surface dyslexia (see Ans *et al.*, 1998; Plaut *et al.*, 1996, for discussion.) Harm and Seidenberg (2004) have recently begun to explore the process of the division of labour in the two routes of the reading model.

Phonological dyslexia, an impairment in reading pseudowords, has been simulated in two main ways. Both methods cause the network to develop an insufficiently general function relating orthography to phonology. The first method reflects the claim we encountered earlier that phonological dyslexia corresponds to phonological representations (and perhaps orthographic representations as well) that have developed with insufficient componentiality (Manis *et al.*, 1993; Plaut *et al.*, 1996). Harm and Seidenberg (1999) implemented this proposal by restricting the initial computational properties of the phonological component of their model (the phonological output layer, its recurrent connections, and its clean-up units). Their manipulations occurred prior to the onset of literacy and included (1) the initial removal of the clean-up units and severing half the recurrent connections between the phonological units, or (2) restricting the size of the weights in the recurrent connections, or (3) making computations within the phonological component more noisy. All of these manipulations resulted in poorer nonword naming, and some of them impacted on exception word reading as well (P1 in Figure 12.1). Brown (1997) also demonstrated that when both orthographic representations and phonological representations are deliberately constructed with reduced componentiality, reduced non-word reading results at the end of training (P2 in Figure 12.1).

AU: Manis et al 1993—not in refs.

The second method of simulating developmental phonological dyslexia seeks to constrain the nature of the computational function that can be learnt between orthography and phonology, rather than the input or output representations themselves. In models with two routes connecting orthography and phonology, manipulations have focused on impairing the direct route, which is better suited to learning relations between individual graphemes and phonemes (Ans *et al*., 1998; Zorzi, Houghton and Butterworth, 1998a). Initial removal of this route compromises the network's ability to learn a simple function relating orthography and phonology applicable to novel words, and so generalization is reduced (Zorzi, Houghton and Butterworth, 1998b) (P3 in Figure 12.1). Brown (1997) used another constraint on the computational function by employing several three-layer networks with reducing numbers of hidden units and comparing them when performance on regular and exception words was matched. Necessarily, this meant that the networks with fewer hidden units had experienced more training. Networks with fewer hidden units were unable to learn a robust function linking orthography and phonology and so showed poor nonword reading (P4 in Figure 12.1).

What can we deduce from this set of results? The variety of network architectures along with the variety of manipulations makes comparisons difficult—would a manipulation that succeeds in one architecture necessarily succeed in another? Before drawing some general conclusions, it will be useful to consider a recent systematic exploration of the range of developmental deficits that can be generated in these types of associative learning models, by Thomas and Karmiloff-Smith (2003a). Although the following simulations were carried out in the domain of inflectional morphology (specifically, English past tense formation), the simulations systematically examined developmental deficits along similar dimensions to those of reading models but in a single base architecture. The simulations therefore allow for convergent evidence on the implications of various sorts of computational deficits for an associative model attempting to acquire a partially regular domain.

In past tense formation, a phonological representation of each verb stem must be related to a phonological representation of the inflected past tense form. Along with each verb stem, word-specific (semantic) information is provided, so that the model can either learn to generate past tense forms based on the meaning or based on emerging regularities between phonological input and output forms. In the English past tense, there is a majority pattern (the add -ed rule, e.g., 'talk' => 'talked'), along with a minority of exceptions or irregulars ('think' => 'thought', 'go' => 'went', 'hit' => 'hit'). Generalization to novel strings should extend the add –ed rule. A developmentally disordered

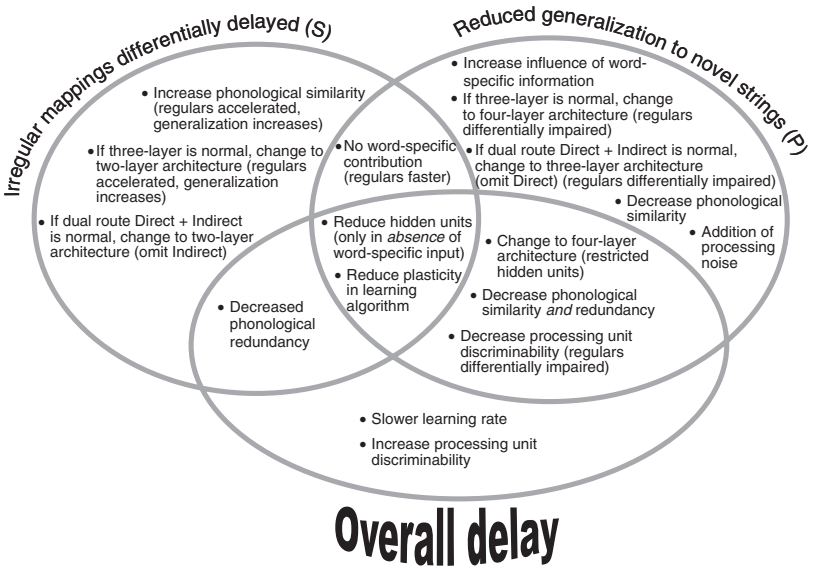

Fig. 12.2 The atypical developmental effects of altering various computational parameters in the startstate of a network that is learning a partially regular cognitive domain analogous to reading, past tense formation (Thomas and Karmiloff-Smith, 2003a). (S) => Delayed irregular acquisition is analogous to surface developmental dyslexia. (P) => Reduced generalization to novel strings is analogous to phonological developmental dyslexia.

model can be assessed on whether it shows (1) a differential impairment in learning *irregular* past tense forms (analogous to surface dyslexia); (2) an impairment in *generalizing* the add –ed rule to novel strings (analogous to phonological dyslexia); and (3) an overall delay in acquiring both regular and exception past tense forms (analogous to delayed reading development). The analogy to the domain of reading breaks down in that, unlike in reading, both input and output representations are in the same modality, and therefore alterations to them must be yoked. In reading, orthography and phonology can in principle vary independently.

Figure 12.2 summarizes the effects of a wide range of startstate manipulations to a normal developmental model along these three dimensions (plotted from Table 1 in Thomas and Karmiloff-Smith, 2003a). The manipulations include changes in the initial network architecture, the initial numbers of hidden units, the processing unit activation function (i.e., the ability of processing units to make fine discriminations), processing noise, the learning algorithm,

the learning rate, the presence or absence of word-specific information, and the structure of the phonological information (where reducing phonological 'similarity' is equivalent to reducing componentiality). Broadly, the manipulations fall under just two groupings: the manipulation either changes the computational problem (specified by input and output representations) or changes the computation power of the learning system (via changes in processing routes, processing resources, activation dynamics, or plasticity).

Figure 12.2 replicates and expands on the pattern of results found in the models of reading acquisition. There are five main conclusions:

1. There are multiple ways of simulating each of the three deficits. Multiple causality of developmental deficits appears to be a strong prediction of both reading and past tense models.
2. Some manipulations produce unique effects on only one dimension, whilst other manipulations produce effects on two or all three dimensions, allowing both 'pure' and 'mixed' patterns of errors.
3. The presence of word-specific information such as meaning, aids the acquisition of exception forms. So too does the presence of a processing route that permits a larger granularity of processing. Initial deficits to either resource impair 'exception' acquisition.
4. The past tense simulations suggest that one finding from the reading models—that reduced numbers of hidden units in the phonological pathway particularly hurt exception mappings—may be an artefact of using models without an implemented semantic route. In the past tense model, the presence of word-specific information allowed the system to overcome the differential effect of hidden unit numbers on exception mappings. In a reading model with a semantic route, particularly under assumptions of division of labour, reduced resources in the orthography to phonology network would likely shift exception processing into the semantic route and allow compensation to recovery.
5. Changes in generalization to novel forms can be disrupted by reduced componentiality/similarity of input representations and of output representations, but also by a range of other factors that alter the mapping function that the system can learn. Note, too, that concentration on a word-specific or semantic route also impairs generalization, since the relevant regularities are not encoded in the dimension of meaning (for reading, they are encoded in a systematic orthography–phonology relationship; for past tense, in a systematic phonology(stem)–phonology(past tense) relationship).

Modelling sensitive periods in reading development

Two other computational models of reading are relevant to the neuroconstructivist principles identified in this book. First, in our discussion of the normal acquisition of reading, we saw the theoretical claim that the onset of literacy serves to restructure pre-existing phonological representations, enforcing a finer level of granularity. Harm and Seidenberg (1999) were able to investigate to the possible nature of this process via simulation work, comparing the structure of phonological representations with and without the influence of orthography. They demonstrated improved segmentation of phonological representations following literacy training, along with sharper representations of rhymes, and a divergence in the phonological representation of rhyming words that had different spellings (e.g., 'hair' and 'bear'). The model illustrates the interactive, activity-dependent nature of a processing system required to link multiple modality-specific representations in driving behaviour.

Second, Harm, McCandliss and Seidenberg (2003) used the Harm and Seidenberg (1999) model under conditions of simulated phonological dyslexia, and evaluated the extent to which there were *sensitive period effects* in remediating the deficit. The dyslexic model was impaired in its initial representations of phonology, as per the Harm and Seidenberg model. The phonological impairment was then alleviated at different points during reading instruction to determine whether this led to improved reading. Harm *et al.*, focused on the sensitivity of this improvement to the time at which the intervention was applied. In this case, intervention was simulated simply by removing the noise/architectural constraints impairing phonological development, as a sort of miracle cure. In a second set of simulations, Harm *et al.*, investigated the efficacy of an existing intervention programme for dyslexic children on a model that retained its atypical constraints. Harm *et al.* (2003) found that in their model, interventions targeting phonological representations had a potential for success in alleviating reading difficulties but only if they were introduced extremely early in learning. Repairing phonological representations once poor learning had become entrenched led to a much-reduced effect. As Harm *et al.*, noted, this replicates data indicating that the degree to which phonological awareness training produces reading improvements critically depends on the time of the intervention with regard to the onset of literacy. Preschoolers show significantly greater benefits than kindergarten or primary school children in their subsequent reading abilities (Bus and Ijzendoorn, 1999). Thus the model demonstrates an increasing *restriction of fate* in the development of its representations, albeit into an atypical state that is insufficient to permit the subsequent normal acquisition of a behaviour based on those representations.

AU: Harm et al 2003—given as in press in refs., which is correct?

AU: Harm et al 2003—given as in press in refs., which is correct?

Why no compensation?

Computational models now allow us to address the 'Why no compensation?' question for developmental dyslexia in mechanistic terms. There are two answers. First, the initial computational constraints of the nascent reading system may militate against compensation. In particular, the implemented models argue that structure–function correspondences exist in the system, so that particular structures (those biased to process particular granularity of mappings, those with access to word-specific information) may be particularly suited to acquiring parts of the reading domain such as single word recognition. If these parts are initially restricted or disconnected, no other part may have the right (or optimized) computational properties to replicate the function. Second, as the empirical data suggests, compensation may indeed occur both in children (Snowling *et al.*, 2003) and in adults (Bruck, 1990). Other structures may attempt to take on the function of the compromised elements. However, they will not be able to compute the relevant functions as efficiently and, moreover, in attempting to compensate, they may interfere with the acquisition of their normal function. A clear example of compensation in computation modelling can be found in the dual-route simulations of Thomas and Karmiloff-Smith (2002a, Figure 6: 746), and in the human case, in the imaging data discussed by Casey, Thomas, and McCandliss (1999).

AU: Casey et al 1999—not in refs.

The crucial point on the prevalence of compensation is that the atypically developing system is defined as disordered *only if its compensatory processes have been unsuccessful*. In empirical studies, disordered participant groups will not include cases where compensation has been successful, unless there is an independent (e.g., genetic) basis on which the disorder can be diagnosed. The at-risk longitudinal studies (Snowling *et al.*, 2003) have indicated that individual variation elsewhere (e.g., while one of the two routes is weaker than normal, the other is relatively *stronger*) may provide opportunities for compensation so that even though the underlying system is atypical, no behavioural disorder is apparent.

The neuroconstructivist principles revisited

Developmental dyslexia illustrates the neuroconstructivist principles in the following way. Reading is a *specialization* of a more general system, driven by exposure to a particular environment and culture. The structure of reading is likely to have *interacted* historically with the constraints of the general system that is recruited, ensuring that scripts are learnable. The location of the emergent reading system in the healthy brain involves *interactions between multiple areas*, determined by which areas are supplied with the appropriate

information and have suitable (often more domain-general) computational constraints required of the task. For instance, it has been argued that the visual word form area arises in a region equipped to resolve the appropriate level of visual detail and compute the appropriate translation invariances. In dyslexia, there may be an absence of emergent specialization, although whether this is a side-effect of a lack of expertise or the cause of subsequent deficits is currently unclear. Different *partial representations* interact across development, for example as illustrated by the restructuring of pre-existing phonological representations at the onset of literacy.

The disorder of developmental dyslexia is continuous rather than all or nothing. It has a genetic component but it is likely that the genetic variations that cause dyslexia also generate anomalies beyond the nascent reading system, both at a neural and cognitive level. However, there is still debate as to the appropriate causal model for these wider anomalies, as well as on the relationship between normal individual variation and atypical development. Poor word–sound processing appears to be a common pathway for multiple causes in the disorder. However, importantly, expression of the disorder depends in part on the nature of the reading environment to which the atypical system is exposed. Phonological dyslexia may represent a poorly developed system for representing spoken speech sounds present prior to literacy across languages, but marked reading deficits only emerge in a subset of the languages that have complex orthography-to-phonology mappings.

Computational modelling has illustrated the importance of evaluating the impact of a clearly specified developmental process in producing the behavioural deficits shown in a disorder. Although there is some way to go to ensure we are capturing the right development process in these models, these models nevertheless represent significant progression from notional extensions of static models to developmental deficits, extensions that fail to specify any mechanisms of change. Models have generated insight into *division of labour (through competition) in multiple component systems*, into *intermodality interactions* across development, into *sensitive periods and progressive restriction of fate*, and into the implications of *multiple causality* in behaviourally defined disorders.

Multiple causality, however, was not unbounded in the models. Broadly, reductions in the similarity or componentiality of phonological/orthographic representations (separately or together) led to reductions in generalization to novel stimuli, while attenuation of the availability of word-specific or semantic information differentially affected exception word performance. But in addition, various changes to the nature of the function that the architecture could compute had overlapping effects. These findings might encourage us to

identify signs of convergence between the modelling of dyslexia and empirical brain imaging data. Imaging work also appears to implicate phonological processing regions and visual processing regions, along with the connectivity between areas representing visual object identity and phonology. However, additional areas have been implicated by the imaging data that have as yet no apparent role in the models (e.g., the cerebellum).

Finally, computational models allowed us to give a clearer answer to the key question in the study of developmental deficits—why, unlike in early-acquired focal brain damage in healthy children, is there no compensation to recovery across development? The answer is that (1) there is compensation, which may produce (perhaps quantitatively, perhaps qualitatively) different cognitive systems; but (2) for those diagnosed with the disorder, compensation is unsuccessful in achieving recovery. The reason it is unsuccessful can be found in a consideration of the computational constraints, typical and atypical, that shape the developmental process when the cognitive system is exposed to a particular—and in this case, cultural specified—environment.

Notes

1. Gene behind dyslexia is discovered. *London Metro*, Tuesday 7 September, 1999.
2. Gene may be a cause of dyslexia. London Metro, Thursday 3 March, 2005.