Connectionist models of development, developmental disorders and individual differences

Michael Thomas

Annette Karmiloff-Smith

Neurocognitive Development Unit, London.

Thomas, M. S. C. & Karmiloff-Smith, A. (2003). 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). American Psychological Association.

Address for correspondence:

Dr. Michael Thomas
School of Psychology
Birkbeck College
University of London
Malet St.
London WC1E 7HX, UK
tel.: +44 (0)20 7631 6207
fax: +44 (0)20 7631 6312
email: m.thomas@psychology.bbk.ac.uk
homepage: http://www.psyc.bbk.ac.uk/people/academic/thomas_m/

Introduction

The computational modelling of cognitive processes offers several advantages. One of the most notable is theory clarification. Verbally specified theories permit the use of vague, ill-defined terms that may mask errors of logic or consistency, errors that often become apparent when formal implementation forces these terms to be clarified. Whereas in the domain of intelligence research, one may refer to a more clever cognitive system as being ‘faster’, an implemented model of that system must specify what ‘speed’ really means. Whereas in the domain of developmental research, one may refer to a more developed cognitive system as containing ‘more complexity’, an implemented model must specify what ‘complexity’ really means. Whereas in the domain of atypical development, one may refer to a disordered cognitive system as having ‘insufficient processing resources’, an implemented model must specify what a ‘processing resource’ really means.

Computer models have recently been applied to each of these domains – individual differences, cognitive development, and atypical development – against a background of pre-existing verbal theories speculating on what cognitive mechanisms might underlie variations in each domain. The aim of this chapter is to examine how computational implementation has forced conceptual clarification of these mechanisms, and in particular, how implementation has shed light on the theoretical relation between the domains. Our discussion focuses on one particular class of widely used model, connectionist networks.

The crux of the issue is as follows. The domains of individual differences, cognitive development, and atypical development each represent a form of <u>cognitive variability</u>: they deal in terms of superior or inferior performance on cognitive tasks. Each computational model contains parameters that alter the system’s performance on the task it is built to address. Therefore, such computational parameters stand as possible mechanistic <u>explanations</u> for variability in performance. Implemented models of individual differences, of cognitive development, and of atypical development have appealed to certain computational parameters to explain superior or inferior performance on cognitive tasks. We can ask firstly, do these models appeal to the same parameters in each case, or different ones? And secondly, what computational role do the parameters play in each model? These two questions can be recast in theoretical terms: do individual differences, cognitive development, and atypical development lie on the same dimension or on different dimensions? And what are the precise computational mechanisms that underlie the dimensions? Our chapter addresses these questions.

In the following sections, we first examine pre-existing theoretical claims that have been made on the relation of the individual differences, cognitive development, and atypical development. Second, we introduce connectionist networks and their component parameters. We then discuss how connectionist networks have been applied to the three domains, in turn cognitive development, atypical development, and individual differences. Third, we compare the three domains, and draw conclusions about the theoretical positions these models embody. Finally, given the aims of this volume, we consider in more depth the form that future computational accounts of individual differences may take, and speculate on whether research might turn up a single 'golden' computational parameter that can explain general intelligence – that is, a parameter that can generate improvements or decrements in performance whatever the cognitive domain.

Pre-existing theoretical claims

(1) How are individual differences and cognitive development related?

First, let us be clear about the target phenomena. By individual differences, we mean the 'general' and 'specific' factors of intelligence. The general factor of intelligence, indexed by Intelligence Quotient (IQ), reflects the fact that individuals tend to show a positive correlation on performance across a range of intellectual tasks. At a given age, the general factor accounts for much of the variability between individuals. In addition to the general factor, there are domain-specific factors such as verbal and spatial ability, which may vary independently within an individual. The exact number of domain-specific abilities is controversial. Individual differences in IQ tend to be relatively stable over time, and IQ in early childhood is predictive of adult IQ level (Hindley & Owen, 1978). This fact suggests that IQ corresponds to some inherent property of the cognitive system. A clue as to the nature of this property might be gained from the fact that performance on elementary cognitive tasks with very low knowledge content correlates with performance on intellectual tasks requiring extensive use of knowledge.

By cognitive development, we mean the phenomenon whereby within an individual, reasoning ability tends to improve with age roughly in parallel across many intellectual domains. Although there may be some mismatch in abilities in different tasks at a given time, by and large children's performance on a wide range of intellectual tasks can be predicted from their age. However, at a certain point in development, children's performance can only be improved to a limited extent by practice and instruction (Siegler, 1978), suggesting that development may not be identical to learning or to the acquisition of more knowledge.

Davis and Anderson (1999) offer a recent, detailed consideration of the theoretical relation of these two forms of cognitive variability. Here we highlight two claims. First, the idea that having a higher IQ is equivalent to having a 'bit more cognitive development' is challenged by the fact that when older children with a lower IQ are matched to younger children with a higher IQ, performance appears qualitatively different. The older children show stronger performance on tasks with a high knowledge component while the younger children show stronger performance on tasks involving abstract reasoning (Spitz, 1982).

Second, several theoretical mechanisms have been proposed to underlie individual differences and cognitive development. In terms of mechanisms that might underlie differences in IQ, several authors have proposed differences in speed of processing among basic cognitive components, on the grounds that speed of response in simple cognitive tasks predicts performance on complex reasoning tasks, and that neurophysiological measures such as latency of average evoked potentials and speed of neural conductivity correlate with IQ (Anderson, 1992, 1999; Eysenck, 1986; Jensen, 1985; Nettelbeck, 1987). Sternberg (1983) has proposed differences in the ability to control and co-ordinate the basic processing mechanisms, rather than in the functioning of the basic components themselves. Finally, Dempster (1991) has proposed differences in the ability to inhibit irrelevant information in lower cognitive processes, since individuals can show large neuroanatomical differences in the frontal lobes, the neural bases of executive function.

In terms of mechanisms that might underlie cognitive development, we once more find speed of processing offered as a factor that may drive improvements in reasoning ability (Case, 1985; Hale, 1990; Kail, 1991; Nettlebeck & Wilson, 1985). Case (1985) suggested that an increase in speed of processing aids development via an effective increasing in short term storage space, allowing more complex concepts to be represented. Halford (1999) proposed that the construction of representations of higher dimensionality or greater complexity is driven by an increase in processing capacity where processing capacity is a measure of the 'cognitive resources' allocated to a task. Lastly, Bjorklund and Harnishfeger (1990) proposed improvements in the ability to inhibit irrelevant information, based on evidence from cognitive tasks and changes in the brain that might reduce cross-talk in neural processing, such as the myelination of neural fibres and the decrease with age in neuronal and synaptic density.

On one hand, then, previous theories relating individual differences to cognitive development proposed that cognitive development is not equivalent to 'more IQ' and thus that development and intelligence are variation on different cognitive dimensions. On the other hand, the lists of hypothetical mechanisms postulated to drive variability in each domain show several overlaps (speed, inhibition), suggesting that development and intelligence could represent variations on the same cognitive dimension(s). There is no current consensus.

(2) Are typical and atypical development qualitatively different?

The relation of typical cognitive development and atypical development could be construed in two ways. Perhaps there are variations in the efficiency of typical cognitive development, whereby atypical development just forms the lower end of the distribution of typical development. This would imply that the two amount to cognitive variation on the same dimension(s) as typical development. On the other hand, one might view atypical development as qualitatively different from normal, as representing a disordered system varying on quite different dimensions.

Current theory holds that individuals with developmental disabilities comprise a combination of these two groups (Hodapp & Zigler, 1999). One group represents the extreme end of the normal distribution of IQ scores in the population (Pike & Plomin, 1996), in which there is no obvious organic damage and individuals frequently exhibit milder levels of impairment. As with typically developing children, individuals within this first group are characterised by relatively even profiles across abilities, albeit at lower overall IQ levels. The second group is more heterogeneous and impairments stem from known organic damage, either of genetic, peri-natal, or early post-natal origin. Although this group shows lower levels of IQ and sometimes severe levels of mental retardation, individual disorders can also demonstrate particularly uneven profiles of specific abilities. For instance, in Williams syndrome, language abilities are often much less impaired than visuo-spatial abilities (see e.g., Mervis, Morris, Bertrand & Robinson, 1999). In Fragile X syndrome, boys can show greater deficits in tasks requiring sequential processing than in those requiring simultaneous processing (Dykens, Hodapp, & Leckman, 1987). And in savant syndrome, individuals with low IQs can nevertheless show exceptional skills within relatively narrowly defined areas such as music, arithmetic, or language (see e.g., Nettelbeck, 1999).

3) General versus specific variation

When we come to examine connectionist approaches to cognitive variability, one distinction will become particularly salient, that between general and specific variation. Theories of individual differences talk about the general factor of intelligence along with multiple independent domain-specific intelligences. Theories of cognitive development stress the apparent general increase in cognitive ability across all domains, but also note the disparities that can emerge between specific domains. Theories of atypical development note that in one group of individuals with developmental disabilities, all cognitive domains are generally depressed, whilst in a second group with apparent organic damage, there can be marked disparities in ability between different specific domains. Any full theory of cognitive variability must address the conditions under which that variability is general across all domains, and when it is specific to particular domains. We will find that connectionist models have generated detailed proposals for specific variability, but thus far have made limited progress on general variability. We now turn to a consideration of these models.

Connectionist models

Connectionist models are computer models loosely based on principles of neural information processing. These models seek to strike a balance between importing basic concepts from neuroscience into explanations of behaviour, while formulating those explanations using the conceptual terminology of cognitive and developmental psychology. (For an introduction to connectionist models, see, for example, Chapter 2 in Elman, Bates, Johnson, Karmiloff-Smith, Parisi & Plunkett, 1996).

Connectionist networks have been widely used to model phenomena in cognitive development because they are essentially learning systems (Thomas & Karmiloff-Smith, 2002). An algorithm is used to modify connection strengths so that the network learns to produce the correct set of input-output mappings by exposure to a training set. By contrast, symbolic, rule-based computational models rarely offer developmental accounts for how relevant knowledge can be acquired, even when such models are able to accurately characterise behaviour in the adult state.

Connectionist models embody a range of constraints or parameters that alter their ability to acquire intelligent behaviours (see Figure 1). This issue is sometimes framed within the nature-nurture debate, in which networks are portrayed as empiricist tabular rasa systems whose knowledge representations are specified purely by their training experience. On closer examination, however, it turns out that in common with all learning systems, connectionist networks contain a set of biases that constrain the way in which they learn. These biases are determined prior to the onset of learning, and include constraints such as the initial architecture of the network (in terms of the number of processing units and the way they are connected), the network dynamics (in terms of how activation flows through the network), the way in which the cognitive domain is encoded within the network (in terms of input and output representations), the learning algorithm used to change the connection weights or architecture of the network, and the regime of training the network will undergo. Only the last of these constraints is derived from the environment; the preceding four are candidates for innate components of the learning system, although in principle these four constraints may themselves be the products of learning.

Decisions about the design of the network directly affect the kinds problem it can learn, how quickly and accurately learning will take place, as well as the final level of performance. To the extent that these networks are valid models of cognitive systems, differences in these constraints or parameters provide us with candidate explanations for the variations found both between individuals and within individuals over time.

To illustrate, networks contain internal processing units that are not specified as input or output units, and are thus available as resources over which the network can develop its own internal knowledge representations (Fig. 1). As we shall see, each of the following claims has been made within the connectionist literature: (1) a network that has more internal processing units is more 'intelligent', i.e., it is able to learn more complex input-output functions; (2) cognitive development can be modelled by networks which recruit extra internal units over time so that more complex ideas can be represented with increasing age; (3) atypical development can be modelled by networks which have too few or too many internal units outside some innately specified normal range.

Figure 1: A typical connectionist architecture: the three-layer feedforward network.

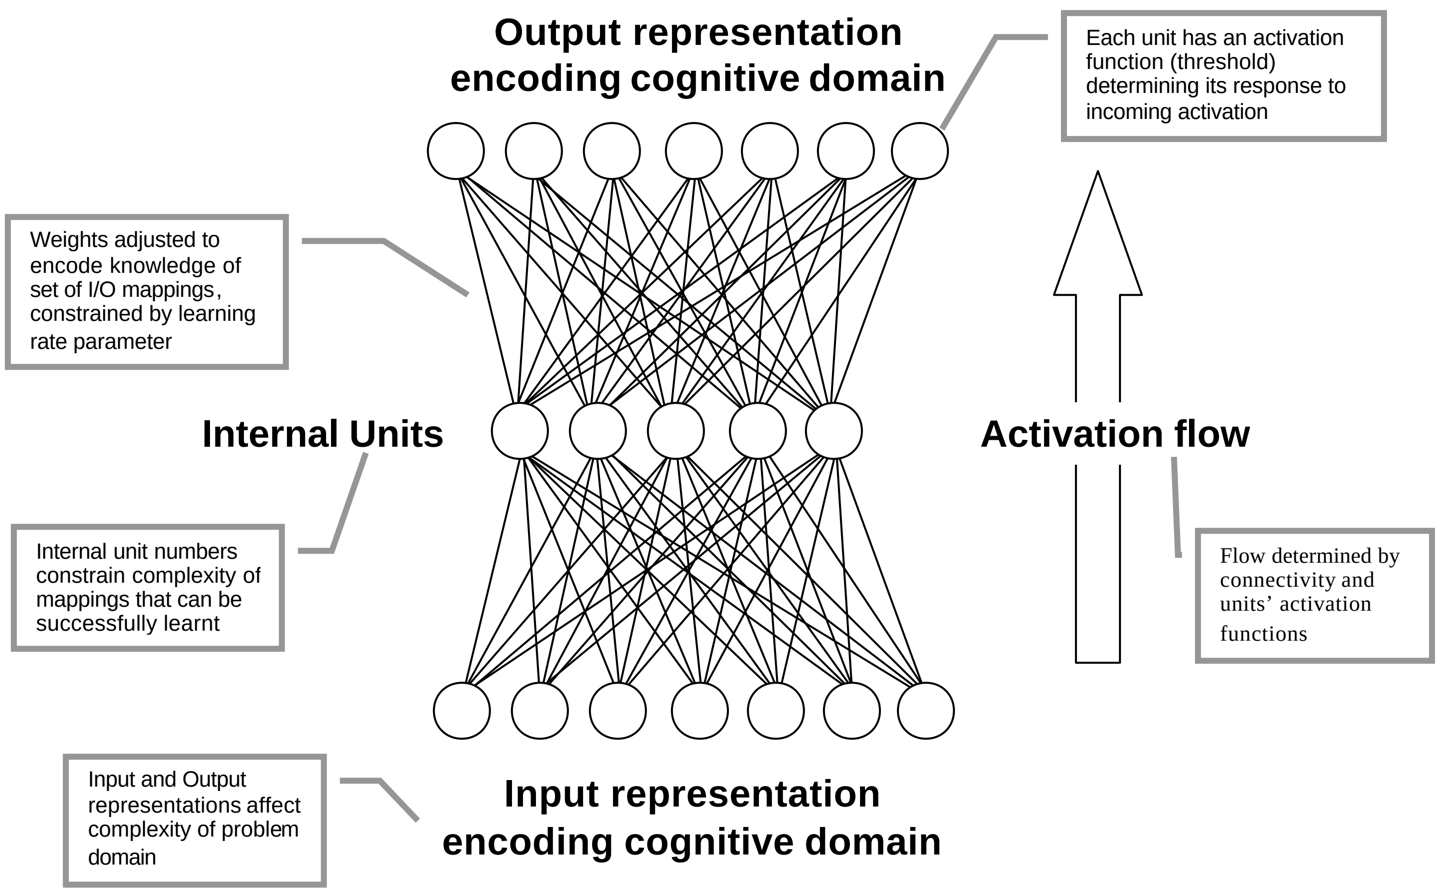

Connectionist models of cognitive development

Connectionist models have been applied to a wide range of developmental phenomena over the last fifteen years. These include categorisation and object-directed behaviour in infants, Piagetian reasoning tasks such as the balance scale problem, seriation, and conservation, and other children's reasoning tasks such as learning the relation between time, distance and velocity, and discrimination shift learning. Within the domain of language acquisition, developmental models have been constructed to investigate the categorisation of speech sounds, the segmentation of the speech stream into words, vocabulary development, the acquisition of inflectional morphology, the acquisition of syntax, and learning to read (see Elman et al, 1996; Thomas & Karmiloff-Smith, in press, for reviews).

Typically these models begin by building an architecture specific to the domain of study, with input and output representations able to encode the relevant cognitive information. For instance, in balance scale problems, a child is presented with a balance scale that has a certain number of weights on each side, positioned at various distances from the fulcrum. The scale is fixed in a level position as the weights are added, and the child must predict which side will drop when it is released. A network model of the development of reasoning in this task has an input representation that encodes both the distances and number of weights placed either side of the fulcrum and an output representation that encodes a prediction of which side of the scale will drop (McClelland, 1989). Other assumptions are built into the model, including the number of internal processing units, the connectivity, and the learning algorithm. Connection strengths are initially randomised. Development is taken to correspond to changes in the connection strengths caused by repeated exposure of the network to the problem domain (in this example, instances of varying numbers of weights placed at different distances from the fulcrum, along with the resulting movement of the balance scale).

This type of model embodies the assumption that development and learning are qualitatively the same kind of thing. However, such a theoretical position is controversial. Connectionists who support this view point out that gradual changes in connection strengths in developmental models are able to simulate behavioural data previously taken to imply qualitatively different stages of development with an associated internal restructuring of representations, for instance, as in the case of McClelland's balance beam model (McClelland, 1989). Moreover, models that rely on changes to connection weights can produce complex developmental trajectories not just through changes to weights in specialised modules but also through changes of weights connecting modules (e.g., Mareschal, Plunkett & Harris, 1999). And indeed, evidence that intense instruction cannot accelerate development does not necessarily rule out a continuum between development and learning; it merely implies that the pace of change must be intrinsically limited.

On the other hand, not all connectionists accept this view: some maintain that learning and development correspond to different dimensions of change. They argue that, on their own, changes to existing network weights are insufficient to capture developmental phenomena, and that changes to other network parameters will be necessary before the models can simulate qualitatively different stages of reasoning. One such change might be a progressive alteration in network architecture, either driven by a fixed maturational timetable or by the dynamics of the learning process itself. For example, the so-called generative connectionist approach has sought to capture cognitive development by progressively adding internal processing units to neural networks during learning (e.g., Mareschal & Shultz, 1996). The number of internal units serves to determine the complexity of the input-output function that the network can learn; given enough internal units, a network can in principle learn an arbitrarily complex function (Cybenko, 1989). In generative connectionist models, then, the distinction between learning and development is reinstated. Changes in weights are taken to correspond to learning, while the addition of internal units (also driven by the learning algorithm) is taken to correspond to the increasing complexity of the representations supported by the cognitive system during development.

A similar proposal comes from work in the modelling of cognitive ageing. Li and Lindenberger (1999) have speculated that, while changes in weights may correspond to learning, cognitive development is driven by changes to a different network parameter, the activation function (or threshold) present in each processing unit. This parameter corresponds to the ability of the network's processing units to discriminate between small differences in input activation. The proposal is that discriminability increases with cognitive development, but decreases during ageing. This idea has the advantage that increases in discriminability may also produce an increase in processing speed and a reduction in interference, thereby offering a way to link together some of the theoretical cognitive mechanisms introduced earlier. However, developmental models have yet to be put forward implementing Li and Lindenberger's suggestion. Moreover, their proposal, and that of the generative connectionists, constitute only two of the network constraints that might be altered to account for a process of cognitive development separate from that of learning. Other parameters remain to be explored.

Connectionist models of development remain limited in at least two respects. First, thus far they have only been applied to specific cognitive domains such as language acquisition or categorisation. Therefore they cannot currently address issues concerning the development of domain-general processing capacities. Second, connectionist networks do not readily represent relational or syntactic information. Therefore at present they cannot be used to evaluate claims such as that of Halford (1999), that cognitive development can be explained as an increase in the ability to represent higher orders of relational complexity. Although network models can embody relational information using synchrony binding (e.g., Hummel & Holyoak, 1997), these models have yet to be extended to the developmental realm. The same is true of rule-based models. By design, they are good at representing relational information (e.g., Jones, Ritter, & Wood, 2000), but they offer poor models of development.

Connectionist models of atypical development

Connectionist models have been extended to account for patterns of atypical development found in developmental disorders. Here researchers take an existing model of a given domain that captures development within the normal population, and attempt to demonstrate that alterations to the initial constraints (in terms of its architecture, dynamics, representations, or learning rule) produce an atypical trajectory of development consistent with the behavioural deficits observed within a disorder (Elman et al, 1996; Karmiloff-Smith, 1998; Oliver, Johnson, Karmiloff-Smith & Pennington, 2000). Since the initial work in this field has been based on existing connectionist models of normal development, it has focused on specific cognitive domains and accepted the position that development can be explained in terms of weight changes within pre-structured networks. In principle, however, the approach is extendible to generative networks, whereby a disorder might constitute a disruption to the process of altering network constraints during development, rather than a disruption to learning by weight change.

A number of specific deficits in development disorders have been investigated. Thomas and Karmiloff-Smith (2003a) have demonstrated that impairments in the acquisition of inflectional morphology in the language of individuals with Williams syndrome (WS) (Thomas et al., 2001) can be explained by changes in the initial phonological and lexical-semantic representations of a model of normal development in this domain. A number of startstate manipulations were examined. Two that were successful in simulating the WS data were the use of initial phonological representations with reduced similarity and redundancy in line with a hypothesis concerning an altered role for phonology in WS, and a deficit in the on-line integration of phonological and lexical-semantic knowledge, also consistent with empirical findings. Hoeffner and McClelland (1993) offered a related account for deficits found in the inflectional morphology of individuals with Specific Language Impairment (SLI). To explain deficits in past tense formation, participle formation and pluralisation, the representations of certain aspects of phonology were weakened in line with a hypothesis concerning perceptual processing deficits in SLI. These aspects carried crucial information about inflectional regularities, and so selectively impaired acquisition of the past tense 'rule'.

Several authors (see e.g., Harm & Seidenberg, 1999) have used models of normal development in reading ability to show how alterations to various initial constraints can produce different types developmental dyslexia. Manipulations to simulate dyslexia have included reducing the number of internal processing units, using a less efficient learning algorithm or a slower rate of weight change, constraining the size of weights in learning, degrading the input and/or output representations, eliminating certain layers of units or connections, adding noise to weight changes, or simply exposing the system to less training.

Cohen (1998) has argued that increases or decreases in the number of internal processing units in categorisation networks can account for various behavioural features in autism. For instance, an increase in internal processing units causes fast initial learning that later regresses. Subsequent categorisation performance is not robust, focusing on details of the original training set rather than abstracting general categories. A reduction in internal processing units on the other hand can lead to a failure to learn in complex domains.

Finally, in the context of Li and Lindenberger (1999)'s proposal that reduced unit discriminability can be used to simulate ageing effects and cognitive development, it is worth mentioning that manipulation of this same parameter, now within the confines of a frontal executive module, has been used to simulate attentional deficits in schizophrenia (Cohen & Servan-Schreiber, 1992). This work exploits a mechanism for attentional selection postulated to operate within domain-specific networks in normal cognition by mediating the activation function of internal processing units (Cohen, Dunbar, & McClelland, 1990).

Connectionist models of developmental disorders have led to theoretical advances in this field in that they have shifted the focus of explanation to the developmental process itself. A previous, widely used approach was based on the static, adult neuropsychological framework (see discussion in Thomas & Karmiloff-Smith, 2003b). This framework sought to explain strengths and weaknesses in developmental disorders in terms of selective preservation of, or damage to, domain-specific modules. However, Karmiloff-Smith (1998) argued that the use of a framework designed to infer the structure of the adult cognitive system from patterns of breakdown in adults with brain damage was inappropriate for the study of an atypical developmental system, since domain-specific processing modules are likely to be an outcome of development rather than a precursor to it. On the other hand, within the neuroconstructivist framework (Karmiloff-Smith, 1998), the strengths and weaknesses found in adults with developmental disorders are the outcome of an atypical developmental process that may include both atypical modularisation and different underlying cognitive processes – even when surface behaviour is apparently preserved. Developmental connectionist models embody this conception, since high-level behavioural deficits in each trained model are the consequence of a developmental process acting on a system with low-level deficiencies in its startstate, rather than a normal fully trained system suffering selective damage to high-level components.

Finally, we should note that (once more) the connectionist modelling work in developmental disorders has focused on specific cognitive domains, with little work examining domain-general effects. At present, such models are silent, therefore, on the more general impairments to IQ found in many developmental disorders.

Connectionist models of individual differences

To date, less connectionist research has been directed at explaining individual differences. However, a number of themes can be discerned in existing work, and researchers are now beginning to specifically target individual differences for explanation via parameter manipulations to their models.

For some time, the connectionist position on individual differences was merely implicit. Such differences related to random variations in initial network conditions or in the training regime to which the network was exposed. As we have seen, the majority of cognitive models begin by constructing a domain-specific network. The connection weights are initially randomised, and the model is then exposed to an environment of randomly ordered training examples, from which it must acquires the given ability. In demonstrating that the performance of a given network does not rely on some particularly advantageous initial set of weights or random order of presentation, it is standard practice to run the model several times using different randomised start states and presentation orders. Performance over several networks is then averaged. The set of networks can be used to demonstrate statistically reliable effects of a task dimension, akin to an experimental design with human participants.

Sometimes implicitly, and sometimes explicitly (e.g., Juola & Plunkett, 1998), variation in the performance of networks with different random start states and training orders has been equated with individual variation among human participants. If this were a full-scale theory, it would apportion individual differences and cognitive development to different though related sources; respectively, the initial weight matrix and subsequent changes to these weights.

However, this source of individual variation is an unlikely candidate to account for a <u>general</u> factor of intelligence, since there is no reason (at least within currently stated modelling assumptions) why an individual should tend to have simultaneously fortuitous random initial weight sets in separate networks dedicated to unrelated cognitive domains. Such an account would need to implicate a single network involved in many domains for such randomness to have a domain-general effect. Note that there are two senses of a domain generality here. One sense is of a module that is a ‘jack-of-all-trades’, participating in the processing of many domains, of which a working memory system is an example. The other sense is of a module with a specific executive or control function, which is then linked to a range of modules involved in processing disparate domains. This is a distinction will be important later.

McLeod, Plunkett, and Rolls (1998) summarised several in-principle proposals for how connectionist networks might account for individual differences. In addition to initial weight states, they suggested in three other sources of variation. The first of these was in the learning rate of the network, in terms of how quickly weights can be changed in response to learning episodes (see Garlick, in press, for a fuller development of this idea). The second was in the number of internal units, and the third was in terms of differences in learning experience, or the training regime to which the network is exposed.

Three recent papers have explicitly aimed to use connectionist models to account for individual differences found in certain language tasks. Plaut and Booth (2000) sought to account for individual differences in semantic priming in a model of word recognition. They appealed to pre-existing differences in perceptual efficiency, and altered the model’s input to reflect the relative effectiveness of low-level perceptual processes (processes not implemented in the model). Plaut (1997) argued that differences in the patterns of acquired dyslexia shown by adults after brain injury could be explained by pre-morbid individual differences in the division of labour in the reading system between two processing routes for naming, one involving semantics, one not. In an illustrative example, Plaut implemented these differences either by varying sources of information external to the model, or by manipulating a weight change parameter in one of the processing routes within the model. This parameter was taken to reflect the degree to which the individual's underlying physiology 'can support large numbers of synapses and hence strong interactions between neurons'. Lastly MacDonald and Christiansen (in press) sought to explain individual differences in linguistic working memory during sentence processing. Importantly, these authors used connectionist models to argue that the very notion of domain general-processing capacity may be faulty, and that any such capacity is inherently linked to (domain-specific) processing and knowledge. (These criticisms were not extended to the other sense of 'domain-generality' which we distinguished above, that of executive control). MacDonald and Christiansen proposed that individual differences in linguistic working memory derive either from different levels of experience with language, or from differences in the accuracy of individuals' phonological representations. Evidence of correlations between individual differences in linguistic and spatial working memory (which would contribute towards a general factor of intelligence), are accounted for either by shared processing components across memory systems, or by similar 'biology and experience underlying language comprehension skills and the skills used in navigation through space'.

Three conclusions can be drawn from the existing connectionist work on individual differences. First, little work has addressed the general factor of intelligence. Accounts of individual variation are restricted to models of specific domains. Second, existing proposals for parameters that might explain domain-specific individual variation include the <u>usual suspects</u>, i.e., the same parameters that we have seen offered to explain variation in cognitive development and atypical development. These include changes in the number of internal units, changes in the learning rate, and changes in initial representations, as well as differences in the level of training determined by some other (model-) external source of variance. Thirdly, although differences in initial randomised weight strengths might offer a distinct candidate to explain (at least domain-specific) individual differences, this hypothesis makes such strong assumptions concerning the homogeneity of all other computational constraints across individuals that it seems unlikely to suffice as a complete theory.[1]

Connectionism and general intelligence

How might one proceed to investigate general intelligence within the connectionist framework? We envisage at least two possible approaches. First, one could search for a one or more parameters in a domain-general processing module that would improve performance in this module. Since the module is involved in many domains, it would explain correlated performance across those domains. Second, one could search for a one or more parameters present in a range of different domain-specific modules that would improve performance in each of these domains. Cross-domain correlations in individual differences would then be explained by yoked changes in these 'golden' parameter, which had a benign computational influence whatever the problem.

With regard to the first approach, the search for a single parameter in a domain-general processing module, recall that we identified two sorts of 'general' system: one applicable to any domain – such as generalised working memory – and the other involved in the control of a range of separate processing modules – such as executive function. However, we have seen that connectionist researchers have argued against notions of generalised working memory, while detailed models of executive function have yet to emerge.

[1] Plaut, for instance, prefers to view individual differences as involving not just variation in numbers of internal units, but also the density of connectivity between layers of units, learning rate, weight decay, strength of input, and processing rate in recurrent networks (pers. comm., May 2000).

A more ready approach, then, might be to consider whether one parameter can improve performance across existing models of disparate cognitive domains. For example, correlations are found between measures of infant preverbal cognition in tasks like novelty preference, and subsequent childhood performance on standardised assessments of intellectual function, as measured by g-loaded tests such as Raven's Progressive Matrices (see Columbo & Frick, 1999). If we had two connectionist models, one of infant novelty preference in categorisation and one of adult performance on a task like the Progressive Matrices test (roughly, analogical reasoning about geometrical patterns), then we would be in a position to look for a single parameter that could improve performance in both domains.

A comparison of this kind is indeed possible. Mareschal, French & Quinn (2000) have proposed a developmental model of infant novelty preference in which a simple network system learns objects presented to it in categories (defined over perceptual features), and then shows a novelty response to new objects outside the category compared to new objects within the category. There is as yet no developmental model of analogical reasoning about geometrical patterns, nor of a system that performs domain-general reasoning. Theorists such as Fodor (2000) have argued that we know so little about domain-general reasoning that modelling its computations would be premature. However, there exists a model of creative analogy formation in adults in the micro-domain of letter strings (Mitchell, 1993; see also French, 1995). This domain is not too dissimilar from that of the Ravens progressive matrices test, in that the model must produce answers to questions of the form "If <u>abd</u> changes to <u>jki</u>, what is the analogous change to <u>mrrjjj</u>?" The model employs a multi-component system in which 'codelets' or micro-agents compete probabilistically to construct a novel representation of the particular analogy in a working memory space, relying on a complex network of domain-specific background knowledge about letters and the formation of letter sequences. Is there a golden parameter that could improve performance in both models, that might make both more ‘intelligent’?

In infant categorisation, increased novelty preference predicts later IQ. The model would generate an increase in novelty preference with a richer featural representation of the input, a faster learning rate (or more training, or more efficient learning), a larger number of internal processing units, greater unit discriminability, and in a recurrent format, faster settling into a stable activation state. The adult analogy model gains its abilities from many domain-specific structures and processes; it is less straightforward to predict parameters that will improve its performance. French (pers. comm., May 2000) suggests the following would produce <u>less</u> intelligent performance: poor perception of the features of the three input strings, noise in the system, a failure to resolve competitive processes, a failure of the micro-agents to focus on less speculative and more consistent interpretations over time, and an a priori preference for the surface attributes of letter strings (such as <u>identity</u>) compared to deeper structural information (such as <u>successor</u>). If we rule out knowledge-based possibilities, on the grounds that performance in elementary cognitive tasks correlates with IQ, the most likely common candidate to improve performance in both these models is the settling of competitive processes, perhaps mediated by fast-cycling recurrent connections, by noise free processing, or by processing units with greater discriminability.

In sum, researchers are beginning to propose parameters of connectionist models that might account for individual differences. All detailed proposals are implicitly aimed towards specific factors of intelligence, in the main because each model captures only one micro-domain. However, as illustrated, the search for computational factors underlying general intelligence is entirely possible within this framework, and such a search may yet link quality of performance with parameters of underlying computation in different domains.

Discussion

We began this chapter by highlighting the clarification that computational implementation brings to verbally specified theories, and we are now in a position to revisit this claim. In the domain of intelligence, one predominant claim was that differences in cognitive 'speed' might underlie variation between individuals of the same age. A computational consideration of this proposal suggested that 'speed' may need to be interpreted in terms of the settling of recurrent circuits if it is to be linked to the quality of processing, and perhaps to the activation function of units if it is to be linked to the quality of representations. In the domain of development, one predominant claim was that increases in 'complexity' drive improved performance over development. A computational consideration suggested that one feasible interpretation of this claim was that complexity corresponds to the number of internal processing units available to form mental representations at a given point in development. Within the generative connectionist framework, such numbers increase with development allowing more sophisticated representations to be learned. In the domain of atypical disorders, one predominant claim was that cognitive systems might be restricted by having 'insufficient processing resources'. A computational consideration in disorders such as autism and dyslexia suggested that one interpretation of this claim is in terms of networks with too few internal processing units to acquire a domain of a given complexity.

In each case, however, work with connectionist networks has also generated new computational parameters as potential explanations of each type of cognitive variability. Moreover, since it is inevitable that connectionist models will become more complex over time, it is likely that new candidate parameters will emerge. Our conclusion here is simply that theoretical ideas benefit immensely by exploration within a computational framework, where the effects of parameter variations can be studied closely in realistic cognitive domains.

What light does the work we have reviewed shed on the possible conceptual relation of individual differences, cognitive development, and atypical development? Table 1 summarises the different computational parameters proposed to account for cognitive variability. It is notable that <u>many of the same</u> parameters have been proposed to account for different forms of variability, for instance, changes to internal unit numbers and changes to unit activation functions. In principle, then, one might take connectionist work as supporting the claim that different forms of cognitive variability lie on the same dimensions, perhaps representing different ranges of parameter values on those dimensions. However, we suspect that some of this overlap represents a historical anomaly, in which investigators have begun to explore each field of variability by altering the first parameters to hand. Currently there is no coherent overall account of cognitive variability within connectionism (indeed, this chapter represents the first systematic comparison). Time will tell whether the dimensions taken to underlie individual differences, cognitive development, and atypical development subsequently diverge.

Table 1. Existing (independent) proposals for parameters within connectionist models that may explain forms of cognitive variability (see text for references of specific proposals).

Learning	Development	Atypical development	Individual variation	Moment to moment variation (attention)	Ageing
Change weights	Change weights		Different initial random weights		Remove weights
	Change architecture (internal units)	Change architecture (internal units)	Change architecture (internal units)		Remove units
	Alter activation function	Alter activation function		Alter activation function	Alter activation function
		Alter learning rate	Alter learning rate		
		Differences in input/output representations	Differences in input/output representations		
		+ others	Is there a		

			GOLDEN domain-general parameter?		

Finally, we return to a consideration of intelligence. Throughout, it has become apparent that the current state of connectionist research focuses on variations within models of specific abilities, with no good account of general cognitive variability. In the previous section, we speculated on the possibility of a Golden parameter that would improve performance whatever the domain. Indeed Garlick (in press) has recently proposed that learning rate – corresponding to neural plasticity – could be just such a parameter. However, Garlick's supporting simulations only demonstrated that on three problems, networks with a faster learning rate required fewer exposures to the training examples to reach ceiling. There was no indication that learning rate alone could produce changes in the quality or abstractness of representations. It remains a possibility that general intelligence is an emergent property of the interaction of many specific components (Detterman, 1986), an idea that future modelling work should address. Our suspicion, however, is that general intelligence may turn out to correspond to variations of a parameter in a domain-general cognitive system that has a 'finger in every pie', such as an executive system. (A similar but more complicated proposal would be that a given parameter varies in all parts of the system but only produces changes in intelligence via the general module. This would explain many more diverse correlations). Our own experience with modelling leads us to think that the search for a golden computational parameter exerting a beneficial effect whatever the domain may be a difficult one, for the majority of parameter changes have effects that depend on the nature of the problem domain.

Some examples illustrate this point. The famous manipulation of internal processing units turns out to be a parameter where increases help in complex domains but hinder in simple domains. The addition of noise is helpful in the acquisition of domains with broad regularities and irrelevant superficial detail, but damaging when categories must be distinguished on the basis of fine detail. A fast settling system may be an advantage in a domain that is internally consistent, but a hindrance in a domain where many soft constraints must be combined to produce a best ‘compromise’ solution. A fast learning system is an advantage for instantly storing memories, but a disadvantage when it comes to extracting prototypical structure across many exemplars.

The precise computational basis of intelligence remains a mystery, but one of the most promising methods to explore this issue is through the application of computational modelling. It is an approach that may yet allow us to link intelligence to many other forms of cognitive variability.

References

Anderson, M. (1992). Intelligence and development: A cognitive theory. Oxford: Blackwell.

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

Bjorklund, D. F. & Harnishfeger, K. K. (1990). The resources construct in cognitive development: Diverse sources of evidence and a theory of inefficient inhibition. Developmental Review, 10, 48-71.

Case, R. (1985). Intellectual development: Birth to adulthood. Orlando, FL: Academic Press.

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.

Cohen, J., Dunbar, K., & McClelland, J. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. Psychological Review, 97, 332-361.

Cohen, J. D. & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behaviour and biology in schizophrenia. Psychological Review, 99, 45-77.

Columbo, J. & Frick, J. (1999). Recent advances and issues in the study of preverbal intelligence, (pp. 43-71) in M. Anderson (1999).

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. Mathematics of Control, Signals, and Systems, 2, 303-314.

Davis, H. & Anderson, M. (1999). Individual differences and development – one dimension or two? (pp. 161-191) in M. Anderson (1999).

Dempster, F. N. (1991). Inhibitory processes: A neglected dimension of intelligence. Intelligence, 15, 157-173.

Detterman, D. K. (1986). Human intelligence is a complex system of separate processes, (pp. 57-61) in R. J. Sternberg & D. K. Detterman (Eds.), What is intelligence: contemporary viewpoints on its nature and definition. Norwood, NJ: Ablex.

Dykens, E. M., Hodapp, R. M. & Leckman, J. F. (1987). Strengths and weaknesses in the intellectual functioning of males with fragile X syndrome. American Journal of Mental Deficiency, 92, 234-236.

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.

Eysenck, H. (1986). The theory of intelligence and the psychophysiology of cognition, (pp. 1-34) in R. J. Sternberg & D. K. Detterman (Eds.), What is intelligence: Contemporary viewpoints on its nature and definition. Norwood, NJ: Ablex.

Fodor, J. A. (2000). The mind doesn't work that way: The scope and limits of computational psychology. MIT Press, Cambridge (Mass).

French, R. M. (1995). The subtlety of sameness: A theory and computer model of analogy-making. MIT Press, Cambridge, MA.

Garlick, D. (in press). Understanding the nature of the general factor of intelligence: The role of individual differences in plasticity as an explanatory mechanism. Psychological Review.

Hale, S. (1990). A global developmental trend in cognitive processing speed. Child Development, 61, 653-663.

Halford, G. (1999). The development of intelligence includes the capacity to process relations of greater complexity, (pp. 193-214) in M. Anderson (1999).

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

Hindley, C. B. & Owen, C. E. (1978). The extent of individual changes in IQ for ages between 6 months and 17 years, in a British longitudinal sample. Journal of Child Psychology and Psychiatry, 19, 329-350.

Hodapp, R. M. & Zigler, E. (1999). Intellectual development and mental retardation – some continuing controversies, (pp. 295-308) in M. Anderson (1999).

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.

Hummel, J. E. & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. Psychological Review, 104, 427-466.

Jensen, A. (1985). Techniques for chronometric study of mental abilities, (pp. 95-99), in C. R. Reynolds and V. L. Wilson (Eds.), Methodology and statistical advances in the study of individual differences. New York: Plenum.

Jones, G., Ritter, F. E., & Wood, D. J. (2000). Using a cognitive architecture to examine what develops. Psychological Science, 11, 1-8.

Juola, P. & Plunkett, K. (1998). Why double dissocations don't mean much. In M. A. Gernsbacher & S. J. Derry (Eds.), Proceedings of the 20th Annual Conference of the Cognitive Science Society (pp. 561-566). Mahwah, NJ.: Lawrence Erlbaum Associates.

Kail, R. (1991). Developmental change in speed of processing during childhood and adolescence. Psychological Bulletin, 109, 490-501.

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

Li, S.-C. & Lindenberger, U. (1999). Cross-level unification: A computational exploration of the link between deterioration of neurotransmitter systems and the dedifferentiation of cognitive abilities in old age, (pp. 103-146) in L.-G. Nilsson & H. Markowitsch (Eds.). Cognitive neuroscience of memory. Toronto: Hogrefe & Huber.

MacDonald, M. C. & Christiansen, M. (in press). Reassessing working memory: A comment on Just & Carpenter (1992) and Waters and Caplan (1996). Psychological Review.

Mareschal, D., French, R. M., & Quinn, P. (2000). A connectionist account of asymmetric category learning in infancy. Developmental Psychology, 36, 635-645.

Mareschal, D., Plunkett, K., & Harris, P. (1999). A computational and neuropsychological account of object-oriented behaviours in infancy. Developmental Science, 2, 306-317.

Mareschal, D. & Shultz, T. (1996). Generative connectionist networks and constructivist cognitive development. Cognitive Development, 11, 571-603.

McClelland, J. L. (1989). Parallel distributed processing: implications for cognition and development, (pp. 9-45) in R. G. M. Morris (Ed.) Parallel distributed processing: Implications for psychology and neurobiology. Oxford: OUP.

McLeod, P., Plunkett, K. & Rolls, E. T. (1998) Introduction to connectionist modelling of cognitive processes. Oxford: OUP.

Mervis, C. B., Morris, C. A., Bertrand, J., & Robinson, B. F. (1999). William Syndrome: Findings from an integrated program of research, (pp. 65-110) in H. Tager-Flusberg (Ed.), Neurodevelopmental disorders. Cambridge, Mass: MIT Press.

Mitchell, M. (1993). Analogy-making as perception: A computer model. Cambridge, MA: MIT Press.

Nettelbeck, T. (1987). Inspection time and intelligence, (pp. 295-346) in P. A. Vernon (Ed.), Speed of information-processing and intelligence. Norwood, NJ: Ablex.

Nettelbeck, T. (1999). Savant syndrome – rhyme without reason, (pp. 247-273) in M. Anderson (1999).

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.

Plaut, D. C. (1997). Structure and function in the lexical system: Insights from distributed models of word reading and lexical decision. Language and Cognitive Processes, 12, 767-808.

Plaut, D. C. & Booth, J. R. (2000). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. Psychological Review, 107, 786-823.

Pike, A., & Plomin, R. (1996). Importance of non-shared environmental factors for childhood and adolescent psychopathology. Journal of the American Academy of Child and Adolescent Psychiatry, 35, 560-570.

Siegler, R. S. (1978). The origins of scientific reasoning, in R. S. Siegler (Ed.), Children's thinking: What develops? Hillsdale, NJ: Lawrence Erlbaum Associates.

Spitz, H. H. (1982). Intellectual extremes, mental age, and the nature of human intelligence. Merrill-Palmer Quarterly, 28, 167-192.

Sternberg, R. J. (1983). Components of human intelligence. Cognition, 15, 1-48.

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. (2003a. Modelling language acquisition in atypical phenotypes. Psychological Review.

Thomas, M. S. C. & Karmiloff-Smith, A. (2003b. Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. Behavioral and Brain Sciences.

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