Modelling Typical and Atypical Cognitive Development:

Computational constraints on mechanisms of change

Michael S. C. Thomas

Annette Karmiloff-Smith

Neurocognitive Development Unit, Institute of Child Health, London

To appear in: Goswami, U. (in press) *Handbook of Childhood Development*. Blackwells Publishers.

Running head: Modelling cognitive development

Address for correspondence:

Dr. Michael Thomas,
Neurocognitive Development Unit,
Institute of Child Health,
30, Guilford Street,
London WC1N 1EH
UK
Email: M.Thomas@ich.ucl.ac.uk
Tel.: 020 7905 2747
Fax: 020 7242 7177

Introduction

Empirical studies of cognitive development usually report the abilities that children display at different ages. The cognitive mechanisms that allow the child to move from one set of abilities to a more complex set remain shrouded in mystery and have given rise to much controversy.

To take a famous example, Piaget characterised cognitive development as a process of acquiring mental representations of increasing complexity. He proposed that the mechanism of change involved a combination of three processes: assimilation, accommodation and equilibration. One process interprets new experience according to existing knowledge (assimilation) while a second adjusts existing knowledge to fit with new experience (accommodation). These first two are local processes, whereas a third (equilibration) is the attempt of the whole system to find global equilibrium after multiple local changes. In this theory, successive stages of cognitive development have greater complexity and representational power than the previous stages (Piaget, 1954). However, Fodor (1980) argued that, in principle, increases in representational power could not be the consequence of a learning mechanism. This is because the achievement of such increases would require the learning mechanism to evaluate information it did not have the power to represent. Put another way, the mechanism would have to determine the truth of theories that it did not have the ability to understand. Fodor concluded that any increases in complexity during cognitive development are necessarily maturational and that learning is merely a process that uses experience to select among subsets of representational primitives already available to the cognitive system at that point in development.

These two theories place radically different amounts of weight on the role of learning in driving cognitive development. Which is the right account? Part of the difficulty in evaluating the relative merits of these kinds of proposals about mechanisms of change is that verbally expressed theories are often vague and ill-defined. What exactly are equilibration or representational power? What do real learning mechanisms look like and what factors affect the way they learn? Computational modelling offers one method to explore questions like these with far more precision. Models provide the opportunity to establish what types of learning system can be successful in acquiring certain competencies, what constraints such systems should include to best make use of the knowledge available to them in the learning environment, and what stages of performance such systems go through before achieving mastery. Computational models provide candidate systems for the mechanisms of change that drive cognitive development.

In this chapter, we examine the use of computational models for studying development from one main perspective. This is the approach that employs connectionist models, also known as artificial neural networks. Although we relate these models to other types of computational modelling, much of the chapter is taken up with considering the range of cognitive developmental phenomena to which connectionist models have so far been applied, both in typical and atypical populations. We start with a very brief introduction to the basic concepts of connectionist modelling and then consider a single model in some detail, that of children's performance in reasoning about balance scale problems. Subsequently we look at models proposed to account for the development of other aspects of reasoning in children, development in infancy, and the acquisition of language. We then pause to examine some of the theoretical issues raised by these models. In the second half, we consider a recent extension of connectionist networks to capture behavioural deficits in developmental disorders.

An introduction to connectionist networks

Connectionist models are computational systems loosely based on principles of neural information processing. It is important to stress that in the current context, they are not intended to be models of neural circuits, but to sit at a higher level of description. Their aim is to incorporate concepts at the cognitive level so that their performance can be evaluated against behavioural data. Connectionist models are relevant to cognitive development for two main reasons. The first relates to biological plausibility. Although there is controversy over whether current connectionist models have abstracted the correct computational primitives from neural circuits, it nevertheless seems likely that computational solutions achieved in these models will be readily implementable in real neural circuits (O'Reilly, 1998). Computational modelling of any sort can be useful in clarifying theories; however, the attempt here is to build models which employ the same style of computation as the brain.

The second reason that connectionist models are relevant to cognitive development is that they are essentially learning systems. A typical model will comprise an initial network structure and a training environment representing the domain to be mastered. With the application of a learning algorithm, the network alters its structure to achieve competency in the domain. In particular, through interacting with the training environment, connectionist networks develop <u>internal representations</u> or knowledge states that permit them to perform the relevant computations. Given that cognition is characterised as the progressive construction and manipulation of mental representations in the brain, together these characteristics make connectionist networks attractive systems to model processes of cognitive development (Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1996; Karmiloff-Smith, 1992; Plunkett & Sinha, 1992).

A brief overview of the components of a connectionist model follows. More detailed introductions can be found in Elman et al. (1996), McLeod, Plunkett and Rolls (1998) and Rumelhart and McClelland (1986). Software is also available with these volumes enabling the reader to build and explore his or her own computational models.

Connectionist systems comprise simple processing units connected together into networks. Each processing unit has associated with it an activation level, analogous to the firing rate of a neuron. Units influence the activity of their neighbours depending on the strength of the connection between the units. Thus, activity on one unit may facilitate or inhibit the activity of neighbouring units. A given unit determines its activation depending on the sum of facilitation or inhibition it receives from adjacent units and on the unit's decision function (functions can vary, but they usually take the form of non-linear threshold). Networks are generally organised into layers of units, the particular configuration of which is referred to as the <u>network architecture</u>. Figure 1 shows a typical network architecture, with three such layers: an input layer, a layer of internal processing units (usually called 'hidden' units), and an output layer. This particular network uses a <u>feedforward</u> design, in that activation only propagates in one direction up through the network, from input to output. The training environment for such a network amounts to a set of pairs of input patterns (defined in terms of the activation levels to be applied to the input units) and the corresponding desired output patterns.

Knowledge is stored in the network in terms of the strengths of the connections between the units. The strength of a connection is often referred to as its 'weight'. We will not use this term here to avoid confusion when we come to examine the balance scale model. Learning amounts to iteratively altering the connection strengths by tiny amounts so that ultimately, for each input pattern, the network produces the correct output pattern. A variety of learning algorithms are available, but many take the form of <u>gradient descent</u>; that is, for a given input, an 'error' term is derived which gives an measure of how close the actual output is to the desired output. The connection strengths are then changed in a way that reduces the size of this disparity (or moves the network 'downhill' in error space). The most common such algorithm is called <u>backpropagation</u> (Rumelhart, Hinton, & Williams, 1986).

Although each processing unit is computationally very simple, networks of units can compute complex functions. Indeed, in theory a three-layer network such as that shown in Figure 1 can in learn any arbitrarily complex relation between a set of input-output pairs so long as it is given sufficient 'hidden' units over which to develop its internal representations (Cybenko, 1989).

Besides the feedforward network, there is a range of other architectures. For instance, some include loops of connections so that activation can cycle round within the network. In such attractor networks, for a given input the network must gradually settle into a stable state that forms the output, such that further cycling produces no change in output activations. Cycling activation can be used to provide the network with a memory of previous inputs. This enables it to process sequences of inputs, in so-called recurrent systems. Finally, networks can be used merely to form concise representations of a given set of input patterns. Here, there is no desired output supplied in during training. Instead, the network self-organises its representations to form a concise description of the input set across a small set of abstract features.

Figure 1. A three-layer feedforward network.

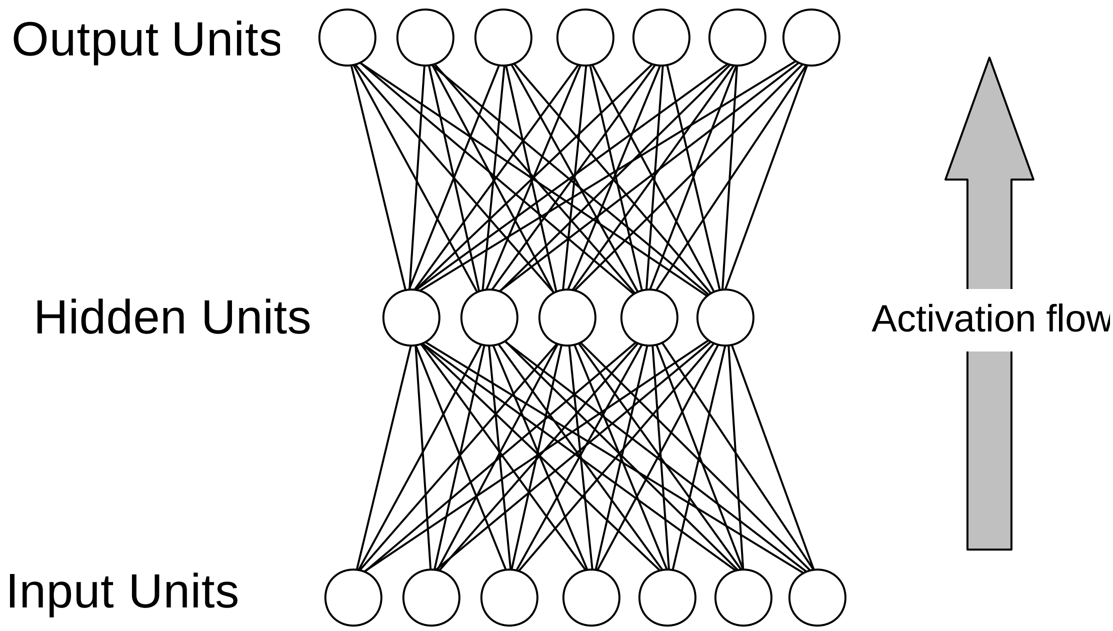

Connectionist models of normal development

An example model: the balance scale task

To illustrate the use of connectionist models for exploring cognitive development, we begin by looking at one model in some detail. The model attempts to capture the development of children's problem solving abilities in balance scale problems. The balance scale was one of a set of problems which Piaget's collaborator, Inhelder, put forward as demonstrating the stages of development through which children pass, each stage representing a more complex understanding (Inhelder & Piaget, 1958). Siegler (1981) demonstrated that children's performance on this task at different stages of development could be characterised by four rules. In the first rule, 4 to 5-year-olds only considered the number of weights on each side of the scale: the side with the greater number will go down. In the second rule, focus is still on the number of weights, but when the weights are equal, the child will then take into account their distance from the fulcrum – the side with the weights further away will drop. The third rule is more sophisticated, always taking weight and distance into account; but if one side has more weight while the other has its weights farther away, the child simply guesses. By age 8, children were generally using rule 2 or 3, and by 12, most children had settled on rule 3. The fourth rule establishes the torque on each side of the scale, with the side with the greater torque predicted to descend. This rule is not reached by everyone: in Siegler's sample, most 20-year-olds continued to use rule 3, although 30% were now using rule 4.

Figure 2a). Structure of McClelland's (1989) balance scale model, with separate channels for processing weight and distance information. Left side of scale: 4 weights on second peg out. Right side of scale: 2 weights on fourth peg out. Scale should balance (both output units half-activated).

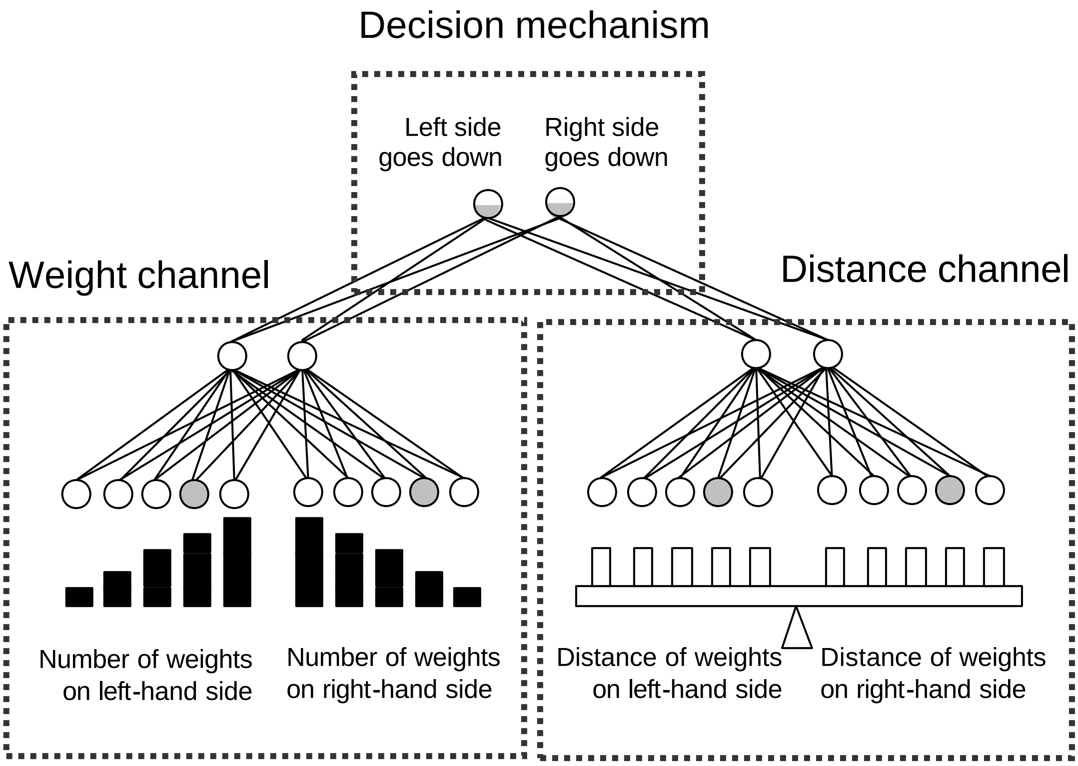

Figure 2b). 'Uninterpreted' structure of the balance scale model – a three-layer network with restricted connectivity between input and hidden layers.

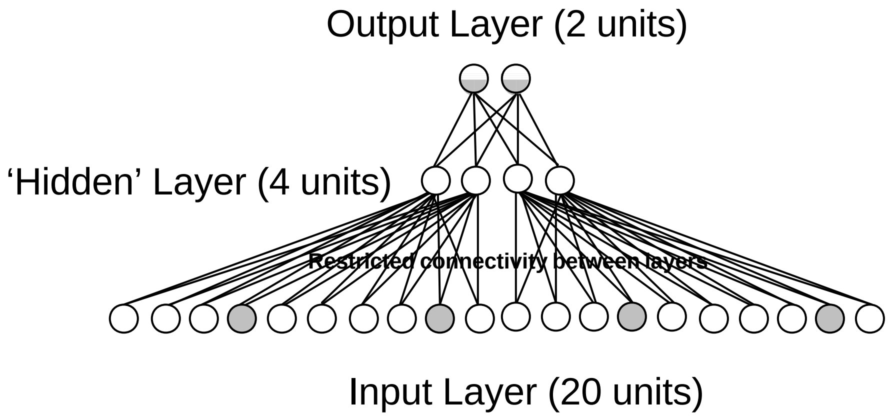

McClelland (1989) proposed a connectionist model of the development of problem solving in the balance scale task (see also McClelland & Jenkins, 1991). The structure of this model is shown in Figure 2a. The input layer is split into two channels, one representing information about the number of weights on each side of the scale, the other representing information about their distance from the fulcrum. Particular configurations can be presented to the model as a pattern of activation over the input units, with 4 units active and 16 units inactive. For each input pattern, the model must determine the side which will drop by turning on one of the two output units and turning off the other. If the scale balances, both output units should be half-activated (as in Fig. 2a). Notice, then, that this domain has been converted into pairs of input-output activation patterns. Note also that the network is oblivious to any information about weights and distances per se. While Figure 2a includes information to help us understand the model's structure in relation to the task, Figure 2b provides a better idea of the model in its uninterpreted form: this is simply a three-layer feedforward network with a restriction on which input units are connected to which hidden units.

Prior to training, the strengths of the connections between the units are randomised within some small range around zero. In this state, the network will perform badly on the balance scale problem because, prior to training, it has no knowledge. However, we will later argue that it does have a sort of knowledge even at this initial stage. For the moment, suffice it to say that it has no knowledge of the content of the problem it is facing (Elman et al., 1996). The model is now exposed to the training set. For each configuration of weights and distances, the network produces an output. This is compared to the correct response, and the connection strengths are altered to bring the network's response closer to the correct output. Connection strengths are altered by only a small amount for each individual problem, and the network is exposed to the set of balance scale problems many times. The consequence is that it converges on a solution where its performance is optimised on the full task. However, in reaching this solution, it goes through phases where it performs well only on subsets of the individual problems. Indeed, McClelland (1989) demonstrated that it passes through the same four stages of performance that the children exhibited and that Siegler characterised with his rules. Importantly, although the four stages appear to be qualitatively different types of behaviour, the model moved between these stages via a single mechanism of continuous and gradual changes to its set of connection strengths.

How does the model simulate this stage-wise pattern of behaviour during development? We have already seen one built-in assumption of the model, that weight and distance information is processed separately. However, a second assumption was built into the model, this time in terms of the training set. The training set was constructed to reflect the supposition that children have a bias to focus more on weight rather than distance. This was achieved by giving the network more training trials where the distance of the weights either side of the fulcrum was equal; only the number of weights varied. As a result of these extra items, the network initially came to rely more on weight than distance as a reliable predictor of the outcome.

Once these two biases are built in, the emergence of the four stages can be explained as follows. Because of the biased training set, the weight channel of the network initially develops more quickly than the distance channel (rule 1). With additional training, the distance channel also begins to develop, but is initially only sufficient to drive the network response when the weight channel does not overrule it (rule 2). With yet more training, the distance channel now starts actively to contribute to the solution, but when weight and distance conflict, the network produces erroneous responses (rule 3). Finally the network finds the appropriate way to combine weight and distance information, effectively discovering the law of torque (rule 4). McLeod, Plunkett, and Rolls (1998) provide a detailed analysis of the connection strengths at each stage within the network. As in the human case, performance of rule 4 is unstable within the model. For the McClelland mode, this is because the law of torque requires a precise calibration of the connection strengths. By contrast, Shultz, Mareschal and Schmidt (1994) have shown that a slightly different model can learn stable rule 4 behaviour, as well as accounting for other empirical effects. Theirs was a generative model, an idea that we will encounter shortly.

In sum, McClelland's model delivers a concrete example of a single mechanism of change in which there is an accumulation of very gradual increases in connection strengths inside the network, but which nevertheless produces behaviour that shifts through stages of qualitatively different performance corresponding to developmental data from children.

Connectionist models of logical development

Seriation is another Piagetian task taken to illustrate distinct stages of transitive reasoning development (Piaget, 1965). Children are initially presented with a random array of sticks of different length, and asked to sort the sticks in increasing length. Four stages of performance have been identified during development: random sorting, sorting of small subgroups, sorting of the entire array but by trial and error, and finally efficient sorting, e.g., by always choosing the smallest stick in the remaining to-be-sorted array.

Mareschal and Shultz (1999) demonstrated how a connectionist network, which once again gradually fine-tuned the strength of its connections, could exhibit a progression of behaviour according to these four stages. Like the balance scale model, Mareschal and Shultz built in an assumption constraining the way in which the model could process the information. They built two sub-networks into the model, one of which was given feedback about which stick to select for moving, while the second was given feedback about where to move the stick. The combination of the development of the 'which' and 'where' sub-networks eventually led to success on the seriation task. Moreover the model predicted that performance would be better when the array was further from its correctly ordered state. The novel prediction was subsequently confirmed by Mareschal and Schultz in testing with young children.

This model differed from the balance beam model in one key respect: during training, the seriation model changed not just the strength of its connections but also its architecture. The seriation model sought to capture the process of development not only by adjusting connection strengths, but also by adding further hidden units when learning faltered. Using an algorithm called cascade correlation (Fahlman & Lebiere, 1990), new hidden units were added not at random but specifically to advance learning when it plateau-ed. Such an approach is referred to as generative connectionism.

Shultz (1998) also used a generative connectionist approach to model children's performance on conservation tasks. These tasks represent a widely used technique for examining children's knowledge of those physical attributes that remain invariant under various sorts of transformation, such as number, weight, and volume (Piaget, 1954, 1965). Shultz constructed a network model of number conservation that was presented with information about two rows of items. The network was told the length and density of each row, with values coded as the activation level of a single input unit for each dimension. This information was presented twice, once describing the rows before a transformation had been performed, and a second time after the transformation had been performed. In addition, the network was given information about which of the two rows had been transformed and about the type of transformation carried out. Two transformations preserved the attribute of number in a row (elongation, compression) while two others altered it (addition, subtraction). The network was trained to indicate whether the rows had the same number of items or if not, which row had the greater number.

As the model altered its connection strengths and added hidden units during training, it simulated a number of the effects exhibited by children. At first, the network failed to conserve number. However, midway through training, the system showed an abrupt shift to conserving number across the relevant transformations. The model also captured the fact that in children, correct conservation judgements emerge for small quantities before they emerge for larger quantities (the problem size effect). Lastly, when the network failed to conserve number appropriately, it tended to choose the longer of the two rows as having more items, just as children do (the length bias effect). The network exhibits the length bias because, in the training set, length tends to be a better predictor of number than density does. Length varies across elongation, compression, subtraction, and addition, while density varies only across elongation and compression (since item spacing is kept constant for subtraction and addition). When the network fails to conserve, it is erroneously employing this 'perceptual' information. The problem size effect in the network is ascribed to the use of continuously valued representations of length and density. With continuous representations, small sizes are easier to discriminate since differences are proportionally larger (e.g. 3 is greater than 2 by 50% but 8 is greater than 7 by only 14%). The continuous coding scheme contrasts with the use of discrete units to code weight and distance values in the balance scale model. In a continuous scheme, 3 (out of a maximum 5) might be represented by turning on a single a unit which normally varies from 0 to 1 to a level of 0.6; in a discrete scheme, 3 might be represented by fully activating the third unit in a row of 5. This contrast illustrates how modellers make different decisions in converting a cognitive domain into a set of activations to be presented to a network.

Lastly, why does the model exhibit an abrupt change from failure to success at this conservation task? Shultz characterises the performance of the network as undergoing a shift from solutions that initially rely on 'perceptual' information, to an increasing focus on the identity of the transformation ('cognitive' information), and finally to a focus on the impact of particular transformations. Shultz attributes the abrupt shift in performance not to incremental changes in connection strengths, but specifically to the addition of extra hidden units. Indeed, some researchers claim that incremental connection changes alone cannot produce discontinuities in development (Raijmakers, van Koten, & Molenaar, 1996) and that changes in architecture are a necessary precondition. However, others argue that the non-linearity of the decision functions within processing units does provide the condition for abrupt changes in performance, even when a network architecture is held constant and its connections gradually changed during training (e.g. Elman, et al., 1996; Plunkett, Sinha, Moller, & Strandsby, 1992).

The models so far described illustrate the broad approach to producing mechanistic computational models of important Piagetian tasks used in exploring childhood development. However, this is not the limit of connectionist models. They have also been use to examine a variety of other developmental phenomena: discrimination shift learning (Sirois & Shultz, 1999), the understanding of the relationship between velocity, time, and distance (Buckingham & Shultz, 1994), the development of concepts (Schyns, 1991), and the development of strategy use in external memory tasks (Bray, Reilly, Villa & Grupe, 1997).

Connectionist models of infant development

Infant behaviour is more closely tied to perceptual-motor skills, and the development of knowledge in infants typically must be explored by indirect means. These include the use of preferential looking and habituation techniques, where infants are found to direct more attention to unfamiliar or unexpected events, as well as the study of search behaviours like reaching for objects. When a ball passes behind an occluder, do infants look to see whether the ball reappears from the other side of the screen? Are they surprised if the ball does not reappear? If the ball remains behind the screen, do the infants reach behind the screen to recover it? The empirical data here suggest a disparity between the different indices of measuring infant knowledge. While infants younger than 7.5 to 9 months fail to reach for the ball when it is hidden behind the occluding screen (Piaget, 1954), infants as young as 3.5 months display surprise in terms of their looking behaviour if the ball does not reappear from behind the screen (e.g. Baillargeon, 1993). If reaching behaviour and looking behaviour disagree, at what age would we want to claim that infants have a mental representation of the hidden object?

Two connectionist models tackled this phenomenon by employing recurrent network architectures. Munakata, McClelland, Johnson, and Siegler (1997) designed a network to predict the motion of objects across the retina. As the system was given increasing experience with the trajectory of moving objects, an internal representation of a temporarily hidden object could eventually be maintained in the memory loop long enough to drive an expectancy that it would re-appear. The model suggests that the concept of the (out-of-view) object is gradually built up through experiences of object trajectories, rather than taking the form of an all-or-nothing understanding of object permanence. Again, performance relies on the fact that knowledge in the network is built up via the gradual strengthening of connections during training. The disparity between expectancy and reaching behaviour is accounted for by showing that the developing representation of the hidden object corresponds to activation patterns which are at first only strong and unambiguous enough to drive expectancy. Further training is required for the representation to be sufficiently robust to drive reaching as well.

Mareschal, Plunkett, and Harris (1999) took a more neurophysiologically motivated approach to the same problem, again using recurrent networks. However, they split processing in the model into an object knowledge module (a 'what' channel) and an object location module (a 'where' channel), based on evidence of such a functional split in the brain (e.g., Milner & Goodale, 1995). A further reaching module had the task of integrating knowledge from both the 'what' and 'where' channels. In this model, the expectancy/reaching disparity was accounted for because expectancy (e.g. looking to where the object might reappear) was driven by one of the lower modules, while reaching required integration of both lower-level modules and so took longer to develop (see Mareschal, 2000, for a review of this research). This explanation predicts that, prior to showing reaching behaviour, infants can generate expectancies based only on information in one of the lower-level modules, relating either to the object's motion or its identity but not both simultaneously. Subsequent empirical evidence has supported this prediction (e.g. Leslie, Xu, Tremoulet, & Scholl, 1998; Wilcox, 1999).

The contribution of models of infant development has been to question the legitimate inferences that should be drawn from the indirect measures of infant knowledge currently in use. Thus far, connectionist models have pointed to the importance of considering graded internal representations, and of considering that certain behaviours may index individual systems while others may require that the cognitive system integrates across separate processing modules and so exhibit developmental lags. While these models have explained high-level conceptual knowledge (categories, representations of hidden objects) in terms of predominantly perceptual information (visual features), they do not imply that there are no such things as high-level concepts, and indeed these models have not provided explanations of the development of conceptual categories. They do however caution against ignoring simpler, low-level explanations of the infant behaviour revealed in current paradigms.

Connectionist models of language development

A great deal of research has employed connectionist models to investigate processes of language development (see Plunkett, 1998 for review). There is insufficient space to review this work here, other than to give an indication of its scope. Models have employed self-organising networks to explore processes of early phonological development in learning to categorise speech sounds (Nakisa & Plunkett, 1998; Schafer & Mareschal, in press). They have used recurrent architectures in models of the development of segmentation of the speech stream into discrete words (Christiansen, Allen, & Seidenberg, 1998) and the acquisition of syntax structure from sequences of words (e.g. Elman, 1993). Generative networks have been employed to model the acquisition of personal pronouns (Shultz, Buckingham, & Oshima-Takane, 1994). Feedforward networks and attractor networks have been used in models of the development of vocabulary (Gasser & Smith, 1998; Plunkett, Sinha, Muller, & Strandsby, 1992), as well as in two of the most heavily researched areas, the development of inflectional morphology, including past tense and pluralisation, (e.g. Daugherty & Seidenberg, 1992; Forrester & Plunkett, 1994; Hahn & Nakisa, in press; MacWhinney & Leinbach, 1991; Plunkett & Juola, 1999; Plunkett & Marchman, 1991, 1996; Plunkett & Nakisa, 1997; Rumelhart & McClelland, 1986), and the development of reading (Harm & Seidenberg, 1999; Plaut, McClelland, Seidenberg, & Patterson, 1996; Zorzi, Houghton, & Butterworth, 1998a).

Theoretical implications

We now turn to consider some of the theoretical issues raised by the use of connectionist networks to model processes of normal cognitive development. The chapter began by pointing out that mechanisms of change in cognitive development were poorly understood, and that verbal theories of such mechanisms, such as Piaget's notions of assimilation and accommodation, were vaguely defined and therefore hard to evaluate. The process of building computational models of development in domains such as conservation and the balance scale task demands that both the task environment and the mechanisms be precisely specified and in turn, allows us to generate candidate realisations of vague theoretical terms.

There have been at least two (contrasting) attempts to use connectionist models of development to offer computational interpretations of Piaget's mechanisms of change. McClelland (1989) suggested that, for networks with a fixed architecture, the activation caused in the network by the presentation of an input might be viewed as assimilation, the interpreting of experience according to existing knowledge. The alteration of connection strengths during learning could then be viewed as accommodation, the adjustment of knowledge structures to fit with new experience. An alternative interpretation of these concepts was offered for generative networks that add additional hidden units during learning. Here assimilation is viewed as the gradual change in connection strengths that occurs during training, while accommodation is viewed as the change in network structure caused when the network adds hidden units to move from a plateau in its learning performance (e.g. Mareschal, 1991; Shultz, 1998).

With concrete definitions of these notions in hand, we can then move on to consider some of the wider theoretical debates. For example, at the beginning of the chapter we saw how Fodor (1980) had downplayed the role of learning in development, arguing that increases in representational power are due to maturation and cannot be a consequence of learning. Learning itself only involves selecting among subsets of pre-existing representational primitives. However, it turns out that Fodor's conception of limited learning is entirely consistent with the learning that goes on in networks with fixed architectures. On the other hand, generative networks that change their structure during learning may indeed be seen as systems that increase their representational power during development as a direct consequence of learning. How can we reconcile the notion that fixed architecture networks are essentially learning systems with the apparently nativist Fodorian notion that learning is a process of merely selecting amongst pre-existing representational primitives? And how is it that generative networks escape this limitation?

To answer these questions, we must consider a description that is often ascribed to connectionist models of development. This is that they are <u>tabula rasa</u> learning systems. If this were true, such models would necessarily fall within a strongly empiricist view of development. However, this characterisation ignores the fact that networks are highly constrained in what kinds of problems they can learn (Karmiloff-Smith, 1992). These constraints are defined by the design decisions that the modeller makes in constructing the network with regard to the initial state of the network, the number of units, layers, connections, and the pattern of connectivity, as well as the way that activation propagates through the network, the learning algorithm, the input and output representations, and the regime of training that the network will undergo. Initial network constraints thus define the full set of representational primitives that the system possesses, and the process of learning selects a particular set. It is in this sense that fixed networks conform to Fodor's nativist notion (Quartz, 1993).

Despite this apparent weakening of the role of learning, the balance scale model illustrates that connectionist networks still manage to give a powerful account of the stage-wise acquisition of problem solving in children. Moreover, this account is consistent with Piaget's constructivist conception of development, as an interaction between pre-existing structures, domain-general learning mechanisms (i.e. principles of neural computation), and the environment. Thus the use of connectionist networks as models of development offers the opportunity of a rapprochement between the apparently divergent views of Piaget and Fodor. Furthermore, these models suggest a view of the nature of innate knowledge. Networks clearly distinguish between knowledge of content, derived through learning, and knowledge based on the constraints that will shape learning. In a connectionist formulation, only the second of these types of knowledge is a realistic candidate for knowledge given to the system in advance of development (Elman et al., 1996).

Although static networks fit Fodor's notion of development, it has been argued that generative networks are examples of systems that can increase their representational power through learning, with the addition of hidden units to improve performance (Quartz, 1993). Thus generative networks seem to be well fitted to model the increase in complexity of cognition during development. However, a caveat must be introduced: to some extent, the claim that generative networks increase their representational power is a matter of frame of reference. From the point of view of the task-specific network, the representational power indeed increases. However, if hidden units are seen as analogous to neurons (or clusters of neurons or cell assemblies or neural circuits), these must have been available in the cognitive system to begin with, albeit not pre-committed to a specific problem. The representational power of the overall system would only change if new hidden units were 'grown from scratch'. On the other hand, if one described a generative network as initially containing a large set of hidden units which it chooses to use only later on in learning (Marcus, 1998a), then the network's representational power would be fixed by this total number of units.

Two other theoretical issues are now briefly considered. The first relates to whether connectionist models are actually sufficiently powerful learning devices to explain the acquisition of the competencies we witness in humans. The second relates to the developmental origin of the network architectures that we see presented for each model. Where do task-specific networks come from?

Learning within any system at all must employ biases to get off the ground (a basic axiom within theories of machine learning; see Mitchell, 1997). The biases contained in connectionist systems correspond to the constraints listed above, and serve just this role of determining what theories the network will consider in learning about a given domain. In a sense, the constraints are the 'helping hand' that the system receives to enable it to succeed at learning. Of course, whether learning is ultimately successful depends not just on this 'helping hand' but also on the complexity of the problem to be learned. For example, a network with too few hidden units may fail to learn a complex set of input-output mappings but be effective in learning a simpler problem. A number of theorists have suggested that certain capacities that humans display could not be learned given the information available in the environment and the constraints typically built into the connectionist networks we have considered (see e.g. Fodor & Pylyshyn, 1988). These capacities usually involve situations characterised by the use of rules, for example in high-level reasoning and language. Rules are hard for connectionist systems to learn because such networks derive their knowledge of content through exposure to associations between inputs and outputs. While networks can extend this knowledge to novel situations, these situations must be similar in some respect to those previously encountered, whereas rules imply generalisation to situations without such similarity. Since humans successfully acquire abilities in rule-based domains such as language, many theorists maintain that the constraints available in network models are insufficient and that these models will require a pre-existing rule-based processing system; the system will initially contain blank rules and then sets the content of these rules through learning (Marcus, 1998b).

Although this view may ultimately be right, three kinds of problems arise. Firstly, current networks can often produce behaviour that looks as if the model were following rules, as if it were simulating a rule-based system (e.g., McClelland, 1989). Secondly, it is not uncontroversially established that humans themselves use rules in cognition. Thirdly, non-connectionist, rule-based computational models of development have great difficulty in accounting for mechanisms of change. By contrast, the very strength of connectionist models is that they are plausible candidates for mechanisms of change in cognitive development. When connectionist networks have been applied to the domain of language, where the case for rule-based representations of syntactic structure appears strongest, recent models of sentence processing have produced an interesting finding. Recurrent connectionist networks that attempt to acquire this rule-based behaviour are only partially successful. However, it turns out that the limitations the networks display conform to similar limitations shown by people when processing sentences (Christiansen & Chater, in press). The current connectionist/rule-based models debate is unlikely to be resolved until either connectionist network models are put forward that convincingly demonstrate the acquisition of apparently rule-based behaviours, or rule-based computational models are put forward that incorporate plausible mechanisms of change.

Lastly in this section, we consider the origin of the architectures proposed to explain performance in each domain. It is a standard assumption in cognitive psychology that the adult cognitive system can be decomposed into specialised components or modules, but some developmentalists claim on neuroconstructivist grounds that high-level modules are more likely the gradual product of development than reflecting the innate structure of the brain (Johnson, 1999; Karmiloff-Smith, 1992, 1998). However, all of the computational models of development we have seen so far focus on explaining only one domain – that is, they assume a module dedicated to learning the computations in a given domain prior to exposure to the training set (see discussion in Karmiloff-Smith, 1992). There is a model for balance scale problems, a model for seriation, and so on, each with input and output representations encoding information specific to its domain. In the broader picture, there is a step missing in connectionist theories of cognitive development. If development itself produces such modules, what is the nature of the developmental process that produces modules? Connectionist work is beginning to address this issue (see Jacobs, 1999, for a review). The basic idea is that although areas of the neocortex may not initially be specialised to particular cognitive domains, they will differ in their computational properties (i.e. in terms of the constraints we discussed earlier). A given set of computational properties equips a given area to be particularly effective in processing a given domain (e.g. recurrent connections would equip a network to process sequences). Areas might thus have <u>domain-relevant</u> rather than <u>domain-specific</u> properties (Karmiloff-Smith, 1998). Their content is not fixed, and in principle, they could be used to acquire other domains for which they are less well suited, although performance would then be sub-optimal. Apart from this computational heterogeneity, areas would also be distinguished by the input and output systems to which they are initially connected. A process of competition between different areas would lead to specialisation and the emergence of modular structure. Thus Jacobs, Jordan, & Barto (1991) showed how specialisation to perform 'what' and 'where' processing in a model of visual object recognition could emerge from a network which at the outset only contained components differing in their computational properties. Work of this nature is a necessary step for connectionist models of development, so that the field can consider some of the more high level aspects of development, such as strategy formation, the interaction between modules (such as between modalities), and phenomena thought to require 'representational re-description' (Karmiloff-Smith, 1992).

Connectionist models of atypical development

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

The first research aim of the field of developmental cognitive neuropsychology has been to characterise the strengths and weakness exhibited in each disorder, i.e. to describe its behavioural phenotype. Temple (1997, p. 5) outlined two further aims. The first is to use developmental disorders to inform and expand our current understanding of normal development. For example, if in a certain disorder, ability A develops in the absence of ability B, one might claim that ability B is not necessary for the development of A. The second aim is to identify what Temple describes as "intact subsystems" within a disorder, which might then be utilised in an educational or remedial context.

Despite the apparent simplicity of these aims, the terminology of "intact subsystems" is in fact controversial, and highlights a current debate in the field. The identification of intact and damaged components of the cognitive system is an approach that originates in research on adults with brain damage. In the field of adult cognitive neuropsychology, patterns of preserved and impaired abilities in adults with different types of brain damage have been used to identify the structure of specialised components within the cognitive system. But Bishop (1997) and Karmiloff-Smith (1997) amongst others have argued that this approach is inappropriate for the study of developmental disorders. Most obviously, using the adult system as a model for a developing infant or child system offers no explanatory role for development in developmental disorders. Use of the adult model in this context relies on two assumptions, both of which are unlikely to be true. The first is that the specialised components found in the adult system are present in the infant system. However, in the previous section, we saw that specialised components are most likely an outcome of development rather than a precursor to it (Karmiloff-Smith, 1992). Second, the adult model needs to assume that if a specialised component is initially damaged in a developmental disorder, the rest of the cognitive system can nevertheless develop normally around it. However, such independence between components does not seem to be a characteristic of normal cognitive development, where early on, interactivity is more typical (Bishop, 1997).

The alternative to the adult brain damage model is to view developmental disorders as the outcome of a long development process occurring in a system in which there are different initial computational constraints (Elman et al., 1996; Karmiloff-Smith, 1998; Oliver, Johnson, Karmiloff-Smith, & Pennington, 2000). Causes are seen not as the failure of high-level specialised cognitive components for, say, language or reasoning about mental states, but in terms of low-level deficits in neural connectivity or the firing properties of particular neurons. These deficits may initially be small, but become exaggerated by the process of development so that marked behavioural deficits are apparent in the adult state. This perspective predicts that atypical systems could show strengths as well as weaknesses, perhaps even demonstrating behaviour at a higher level than typically developing individuals. Even where behaviour appears normal in a given developmental disorder, it may have atypical cognitive processes underlying it (Karmiloff-Smith, 1998).

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

Connectionist models of behavioural phenotypes

Perhaps the largest body of work in this area has been dedicated to the investigation of the possible computational causes of developmental dyslexia. A number of connectionist models have attempted to simulate the reading process in adults, a process characterised in this context as learning to map between representations of the orthographic and phonological properties of word forms, and their corresponding meanings. The general computational framework postulates that hidden units mediate the mappings between these three sources of information, and that processing is both bottom-up and top-down (Seidenberg & McClelland, 1989). In the usual case, however, only a portion of this framework is implemented within a working model. While few current models constitute serious attempts to capture the developmental processes of early reading acquisition (see Zorzi et al, 1998a, for an exception), these models do commence with randomised connection strengths and use training on large word sets to acquire the adult processing structures. On this basis, several attempts have been made to alter initial constraints in these models such that at the end of training, the network exhibits the behavioural features of dyslexia. The target is to capture two particular clusters of deficits. In phonological dyslexia, children and adults predominantly show difficulties in reading novel words. In surface dyslexia, there is a difficulty in reading words whose pronunciations form exceptions to the usual way letters map onto sounds.

With regard to surface dyslexia, an impairment in reading exception words is simulated by altering any initial constraints that reduce the general ability of the network to learn. Exception words will be 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 the network mapping between orthography and phonology (Bullinaria, 1997; Harm & Seidenberg, 1999; Plaut et al., 1996; Seidenberg & McClelland, 1989; Zorzi, Houghton & Butterworth, 1998b), a less efficient learning algorithm (Bullinaria, 1997), less training (Harm & Seidenberg, 1999), and a slower learning rate (Harm & Seidenberg, 1999).

Phonological dyslexia represents a case of developing an insufficiently general function relating orthography to phonology. This could derive from input and output representations that themselves are insufficiently general (Brown, 1997; Plaut et al., 1996; Seidenberg & McClelland, 1989), or network constraints that prevent a general function being learnt even when given appropriate initial representations (Harm & Seidenberg, 1999; Zorzi et al., 1998a). Since reading assumes a large pre-existing spoken vocabulary, some researchers have focused on the development of phonological representations in isolation, following the hypothesis that problems with phonology precede the attempt to relate visual word forms to pronunciations. For example, Harm and Seidenberg (1999) separately manipulated weight change algorithms and architectural constraints in a model of the development of phonology, to explore what initial alterations would generate insufficiently general representations at the end of training.

It is notable that several manipulations have been proposed to simulate each form of dyslexia. To the extent that these manipulations are mutually successful, one might infer that there are many ways to produce the same behavioural deficits. This is a point to which we will return shortly. In terms of dyslexia, the next research step will be to determine which of these manipulations are computationally equivalent accounts and which are sufficiently distinct to generate different empirically testable predictions.

In Specific Language Impairment (SLI), there is a serious limitation in language ability without associated impairments in hearing, low IQ, or neurological damage (Leonard, 1998). It has been proposed that these individuals show relatively greater deficits in using grammatical rules than in accessing lexical items. Problems with inflectional morphology are often cited as an example. For instance, in English past tense formation, individuals with SLI fail to show the usual advantage of regular past tense formation (talk-talked) over irregular (creep-crept), and fail to inflect the root form in many cases, producing 'unmarking' errors (e.g., van der Lely & Ullman, 2001). Hoeffner (1992) constructed a connectionist model of inflectional morphology designed to learn the mappings between the meanings of verbs and their phonological forms under a variety of inflections, including past tense, third person –s and progressive –ing suffixes (e.g. [jump] => jump, jumped, jumps, jumping). Hoeffner and McClelland (1993) then altered the initial constraints under which the model was trained in order to simulate SLI. Their aim was to test the hypothesis that SLI is associated with impairments in the processing of speech which affect the development of phonological representations and in turn, cause knock-on effects in learning morphology and syntax (see Leonard, 1998). On the grounds that children with SLI have difficulty in processing low-phonetic substance inflections in English, Hoeffner and McClelland systematically degraded the phonological representations of their model such that the network's ability to represent word final stops and fricatives (including /t/, /d/, and /s/) was particularly impaired.

The model showed slower and more error-prone learning, with differential deficits across the inflection types. Its final performance displayed a dramatic increase in the number of errors where the verb stem was unmarked, as well as difficulty applying the past tense rule to verbs and, to a lesser extent, in producing irregular past tense forms – all characteristics of SLI. Moreover, just as in SLI, the model showed an impairment on morphemic phonemes (e.g. the final /d/ in died) but not phonologically identical phonemes which were non-morphemic (e.g. the final /d/ in need). This is because during training, 'die' was often presented in contexts other than with the final /d/, such as in 'dies' and 'dying'. On the other hand, since neither 'nees' nor 'neeing' existed in the training corpus, 'nee' was always presented in the context of a final /d/. 'Need' was therefore learned as a single output, while 'died' was learned as a stem with an optional (and vulnerable) affix. This result is important because it establishes the viability of a perceptual deficit account that can preferentially target morphemic phonemes.

Despite some impressive features, this model has limitations. Ullman and Gopnik (1999) point out that the model was unable to capture the low performance on irregular as well as regular verbs reported by themselves and by van der Lely and Ullman (2001). Furthermore, the perceptual deficit account of SLI on which it is predicated remains controversial.

Autism represents a behavioural phenotype with a much wider range of impairments. The disorder is characterised by a central trial of deficits in social interaction, communication and imagination. In addition, there are other associated features, including a restricted repertoire of interests, an obsessive desire for sameness, savant abilities, excellent rote memory, a preoccupation with parts of objects, improved perceptual discrimination, and an impaired ability to form abstraction or generalise knowledge to new situations (see Happé, 1994, for a review).

Cohen (1994) suggested that simple categorisation networks could capture the fact that in some cases, children with autism have trouble acquiring simple discriminations and attend to a restricted range of stimuli, while in others, children have good discrimination and indeed very good memory but seem to rely on representing too many unique details of stimuli. Cohen concluded that evidence from neuropathological investigations of the brains of affected individuals was suggestive of abnormal wiring patterns in various brain regions. In comparison with the normal brain, the structural deficits were consistent with too few neurons in some areas, such as the cerebellum, and too many neurons in other areas, such as the amygdala and hippocampus. Cohen showed that simple classification networks with too few hidden units showed a failure to learn, while those with a surfeit of hidden units showed very fast learning, but subsequently generalisation became poor, and the network increasingly responded according to particular details of the training set. Interestingly in this case, neuropathological evidence was used to motivate alterations to initial constraints in a cognitive model, such that the performance of the network made contact with behavioural deficits of the disorder (see Mareschal & Thomas, in press, for detailed discussion of this model).

Gustafsson (1997) has argued that alterations to the dynamics of learning in low-level sensory features maps could also account for differences in perceptual discrimination in autism. Gustafsson proposed that the relevant atypical computational constraint would be the level of inhibition existing between units in self-organising feature maps. In these maps, units compete with each other to represent particular aspects of the input. If the competition is too fierce, the argument goes, units will end up defending too small a territory – that is, they will come to represent too fine a level of detail in the sensory input to support robust categorisation, although perceptual discrimination will be facilitated. This idea has yet to be tested in a direct implementation, although Oliver et al. (2000) have explored related ideas (see Thomas, 2000).

O'Loughlin and Thagard (2000) have proposed that a similar idea applied to a very much higher level (and less developmental) model may also account for difficulties in theory-of-mind reasoning in autism. Their model exhibited a deficit in reasoning about false beliefs, produced by high levels of inhibition between representations of concepts in a hand-wired interactive network. The system reasons by settling into a stable activation state. Cycling activation is constrained by the inhibitory and facilitatory connections that represent the consistency between beliefs, each belief being represented by a single unit. With heightened inhibition, however, the network falls into a stable state before it has time to integrate all aspects of its knowledge. In particular, there is insufficient time for knowledge about false beliefs to override information about the state of the world directly fed from the perceptual system. The attempt to link low- and high-level models of autistic characteristics via alterations to a similar computation constraint is novel. One might question the validity of attempting to link such disparate levels of description, and the theory-of-mind model certainly sheds much developmental and biological validity. However, together these proposals represent one of the few exceptions to the current trend in this field of addressing the behavioural deficits of a disorder in isolation.

Modelling a developmental disorder with clear genetic cause: Williams syndrome

Williams syndrome (WS) is a rare neurodevelopmental disorder, caused by a micro-deletion on one copy of chromosome 7 (Tassabehji et al., 1999). It results in specific physical, cognitive, and behavioural abnormalities (Karmiloff-Smith, 1998). The syndrome has been of particular interest to cognitive scientists because individuals with WS exhibit an uneven cognitive-linguistic profile, together with mental retardation (Howlin, Davies, & Udwin, 1998). Full IQ scores, typically between 50 and 70, mask differences in specific cognitive abilities: individuals with WS frequently display relatively good verbal abilities alongside deficient visuospatial abilities (e.g. difficulties in constructing patterns, drawing, etc.). While this is the most salient disparity, there are others. People with WS often perform within the normal range on certain standardised tests for face recognition (Bellugi, Wang, & Jernigan, 1994; Udwin & Yule, 1991), and show relatively good performance on theory-of-mind tasks (Karmiloff-Smith, Klima, Bellugi, Grant, & Baron-Cohen, 1995). By contrast, they exhibit difficulties in numerical cognition (Karmiloff-Smith et al., 1995), and in problem solving and planning (Bellugi, et al., 1994).

The dissociation of cognitive abilities in WS has led to the use of this syndrome to support arguments concerning the independence of certain cognitive abilities during development, in particular the developmental independence of general cognition and language. However, it has also been used to attempt a fine-scaled fractionation of the language system itself. Thus, Pinker (1994, 1999) and Clahsen and Almazan (1998) have claimed that individuals with WS have intact mental representations of grammatical knowledge but an impairment to the system which stores knowledge about individual words. The genetic nature of the syndrome then leads to the claim that dissociations found in the language abilities of individuals with WS can serve to reveal the innate structure of the language system – for instance, that the distinct processing of grammatical rules versus word knowledge is built into the cognitive system prior to birth.

The experimental evidence put forward to support this picture once again comes from the acquisition of the English past tense, where it has been claimed that individuals with WS show a specific difficulty in producing irregularly inflected past tense forms (Clahsen & Almazan, 1998). Theoretical arguments are based on a particular (verbal) model of the acquisition of past tense formation (Pinker, 1994). Pinker proposed that the acquisition of the past tense involves two mechanisms, one that learns the grammatical rule (add –ed to the verb stem), the other that learns about individual words that are exceptions to the rule. (This is a good example of a theoretical approach wishing to ascribe an innate rule-processing mechanism to the cognitive system to permit it to learn rule-based behaviour). If separate mechanisms indeed underlie formation of regular and irregular past tense forms, then it is assumed that difficulties in inflecting one of these forms can be taken as revealing impairments to the respective mechanism. In WS, then, the evidence would imply a deficit to the mechanism processing irregulars. We saw earlier the claim that individuals with SLI have difficulties forming <u>regular</u> past tenses. By the same logic, SLI would imply a disorder with a deficit to the regular rule mechanism. Taken together, Pinker (1999) has argued that these two disorders represent a 'genetic double dissociation' of two mechanisms in the language system (p. 262).

Three points are of interest here. Firstly, Pinker's theory is not the only one that seeks to explain the acquisition of inflectional morphology. A number of neural network models have been proposed seeking to explain how this partially regular domain might be acquired (e.g., Plunkett & Marchman, 1991; Plunkett & Juola, 1999). These models have demonstrated that a neural network learning the relationship between phonological representations of verb stems and past tense forms, can successfully acquire both regular and irregular forms in the same network, as well as extend the 'add –ed' rule to novel exemplars (see Pinker, 1999, Thomas & Karmiloff-Smith, 2001a, for discussions). These models make no <u>a priori</u> division in their architecture between processing structures for regulars and irregulars. In contrast to the verbal theory offered by Pinker, these connectionist systems are working computational models of typical development, evaluated both against the type and quantities of errors they make during learning and against their success in acquiring the domain. We will shortly see this alternative model used to explore atypical language development in WS.

Secondly, the argument that deficits in WS relate to damage to particular processing structures in a model of adult performance represents a classic example of an attempt to extend a brain damage approach to a developmental disorder. As an illustration, here is the claim made by Clahsen and Almazan on the basis of poor WS performance on irregular past tense formations: '[the] computational [rule-based] system for language is <u>selectively spared</u>, yielding excellent performance on syntactic tasks and on regular inflection, whereas the lexical system and/or its access mechanism required for irregular inflection are <u>impaired</u>' (1998, p. 193, italics added). The explanation of the behavioural deficit is couched in terms of damage to an adult system – an impairment to lexical memory – alongside other components which are claimed to be spared or intact. The proposal contains no role for development, and the 'intact' component is assumed to have reached its normal endstate. This explanatory framework becomes obvious when Pinker directly compares the behaviour of individuals with Williams syndrome to past tense deficits shown by patients with Alzheimer's disease and aphasia (Pinker, 1994).

Thirdly, the initial data collected to support this claim had limitations, both in the small number of participants with WS involved in the study and in the nature of the control data against which their performance was compared. Subsequent to these initial findings, the largest study to date (Thomas, Grant, Barham, Gsödl, Laing, Lakusta, Tyler, Grice, Paterson & Karmiloff-Smith, 2001) suggested that the performance of individuals with WS is in fact best captured by three characteristics. Firstly, children and adults with WS are delayed in their production of past tense forms, showing a level of accuracy demonstrated by much younger typically developing children. This is consistent with other findings that language development in WS is often delayed (e.g., Harris, Bellugi, Bates, Jones & Rossen, 1997). Secondly, while individuals with WS do show poorer performance on exception verbs than regular verbs, this appears to be in step with their delayed performance, since younger children also find exception verbs harder than regular verbs. There is no specific deficit for irregular verbs. Thirdly, individuals with WS are significantly less willing to generalise what they know about existing verbs to novel verbs, for instance in extending the regular rule (crog-crogged).

Thomas and Karmiloff-Smith (2001a) set out to explore whether alterations to the initial constraints of a connectionist model of past tense development could account for these three features of the WS data. The past tense network mapped from verb stem to past tense form in the presence of semantic information. Various claims have been made that there are in fact subtle deficits in the language system of individuals with Williams syndrome. These include the proposals that their phonological representations may be atypical and perhaps rely on sensitive auditory processing (Karmiloff-Smith, Grant, Berthoud, Davies, Howlin & Udwin, 1997; Majerus, Palmisano, van der Linden, Barisnikov & Poncelet, 2001; Neville, Mills, & Bellugi, 1994;), that their semantic representations may be atypical (Rossen, Klima, Bellugi, Bihrle, & Jones, 1996; Temple, Almazan, & Sherwood, in press), or that semantic information about words may be poorly integrated with phonology (Frawley, in press; Karmiloff-Smith et al., 1998). In order to explore the viability of these different accounts to explain the pattern of performance in the past tense task, Thomas and Karmiloff-Smith altered the initial constraints of the network model to implement each type of deficit.

They found the following results. First, a manipulation of the phonological representations that reduced their similarity and redundancy was sufficient to reproduce the delay for regular and irregular past tense forms, as well as the reduction in generalisation. Second, the pattern could also be produced when noise was added to the information coming from the semantic system during the acquisition of the past tense. Third, elimination or weakening of the semantic contribution produced a pattern inconsistent with the WS data, including a selective delay for irregular verbs and no reduction in generalisation. Lastly, slowed learning failed to produce a reduction in generalisation, suggesting that delayed development alone was insufficient to explain WS performance, and atypical computational constraints are involved. This modelling work was therefore able to test the viability of several competing hypotheses on the causes of language impairments in Williams syndrome. Manipulations to phonology or to the integration of phonology and semantics were able to simulate the past tense data; manipulations to semantics alone or delayed development were not.

What if the WS data had shown a selective deficit on irregular verbs – could the model have shown this pattern? In addition impoverishing semantic information, performance on irregular verbs could be preferentially delayed either by employing a two-layer network which restricted the complexity of the function that the network could learn between verb stems and past tense forms, or by a calibrated reduction in the plasticity of the learning algorithm. However, none of these three manipulations resulted in an impairment to irregular verbs at the end of training. Endstate deficits on irregular verbs only emerged if training was terminated earlier than usual or if the network was forced to learn at a slower rate, so that by the end of training, regular verbs had reached ceiling but irregular verbs had not. In this latter case, a selective deficit for irregular verbs would then be apparent.

Although not necessarily relevant to Williams syndrome, the potential of the past tense model to produce a selective developmental deficit in irregular past tense formation leads to an important theoretical point. When Thomas and Karmiloff-Smith's findings are taken in combination with Hoeffner and McClelland's (1993) model of past tense formation in SLI, it is evident that alterations in initial constraints are sufficient to produce either selective impairments in regular or irregular past tense formation at endstate. Compared to Pinker and Clahsen and Almazan's static explanation based on selective damage to an adult model of the past tense system, these computational developmental models give a more plausible explanation of deficits in developmental disorders. In addition, the two models imply that developmental double dissociations should not automatically be taken to reveal structure within the cognitive system: neither of these models employed a structural distinction between components for processing regular and irregular verbs, yet regular and irregular verb performance could be selectively dissociated by different changes to the initial network constraints.

Modelling implications for atypical development

We will use the Thomas and Karmiloff-Smith model to focus on a number of theoretical issues. The first issue addresses how easy it is to produce a given pattern of developmental impairments in these models. A number of connectionist models of developmental disorders have simply demonstrated that manipulating one network constraint is sufficient to capture some target atypical data. But how many other network constraints did the researchers test? What if many possible manipulations applied to the model at the start of training successfully simulate the target data? Then the simulations would be consistent with many different theoretical accounts of the impairment. The model would fail to usefully constrain the theory (unless, of course, the theory was that a given impairment had many causes). When a model with a single altered constraint simulates a pattern of atypical data, how can we be confident that we have the right explanation for the emergence of a behavioural deficit?

In addition to their theoretically-driven manipulations, Thomas and Karmiloff-Smith (2001a) systematically explored a range of other manipulations to the initial state of the model, including varying the number of hidden units, varying the architecture, and varying the learning algorithm, adding noise to processing, and altering the threshold function in the processing units. When they compared the results, they found that the three features of the WS data were in fact generated by very few manipulations. However, interestingly, many of the manipulations produced one or two of the three features. A single endstate deficit might be the result of a number of different initial network manipulations, whereas patterns involving several features were harder to come by. One might take this result as offering a cautionary note for disorders that are defined on the basis of narrow behavioural impairments (e.g. 'grammatical' SLI; van der Lely, 1997). The modelling work suggests a high risk that such disorder groups will contain individuals for whom the computational cause of the impairment is different. How can we avoid this? Simulations generated a possible solution. Groups of individuals who share the same underlying cause are likely to show smaller levels of variance across other related behavioural measures, and across longitudinal testing, compared to groups of heterogeneous cause (see Thomas & Karmiloff-Smith, 2001b).

We have encountered the theoretical debate that adult brain damage models are inappropriate for explaining behavioural deficits in developmental disorders. Here again, developmental connectionist models can help us evaluate this claim. Connectionist models of adult performance have been widely used to simulate cognitive breakdown under brain damage (see e.g., Reggia, Ruppin, & Berndt, 1996). Now we can straightforwardly compare the deficits in performance of a network model that experiences the same damage prior to training (as in a developmental disorder) with one that experiences it at the end of training (as in adult brain damage). Are the behavioural deficits the same in each case? Thomas & Karmiloff-Smith (2001c) carried out this comparison for several forms of damage in the past tense network: the elimination of network connections, the addition of noise to unit activations, and alterations to the ability of units to discriminate activation levels (i.e., the sharpness of their activation thresholds).

The results showed a highly complex pattern relating adult (endstate) and developmental (startstate) deficits. Sometimes damage at the two points produced the same effect (eliminating connections) but to different degrees (removing connections in the endstate was far more damaging, since the network could no longer reorganise around the damage). Sometimes deficits appeared in one case but not in the other (higher discrimination produced deficits only when applied to the endstate; noise produced deficits only when applied to the startstate). Sometimes the effect of a manipulation was selective (higher discrimination only impaired performance on exception verbs), whereas sometimes it was global (noise caused a deficit for regulars, exceptions, and in generalisation). Given that the startstate and the endstate are separated by a dynamic developmental process, it is perhaps unsurprising to find in this model that the relationship between them is so complex. This modelling work does, however, tend to support the contention that adult deficits are unlikely to serve as a reliable analogy to developmental deficits.

Finally, in this chapter we have encountered modelling work that seeks to explain two types of variation: the variation in cognitive performance as individuals get older, and the variation between individuals who are developing typically compared to those who are developing atypically. There is a third type of variation, however, that of individual differences or intelligence. There is insufficient space here to describe the initial connectionist work on modelling intelligence. However, it is worth mentioning that it is an open question whether each type of variation should be explained in these cognitive models by manipulating similar computational parameters, or whether each type of variation should be explained by appeal to different parameters. Indeed, there is a further debate about whether learning and development should themselves be ascribed to different mechanisms of change. Thomas and Karmiloff-Smith (in press) recently compared connectionist models of the three types of variation. They concluded that thus far, similar manipulations have been used to model each type of variation (a popular candidate being alterations in the number of hidden units). However, this similarity may be an artefact of the infancy of the respective fields, and explanations for sources of variation may well diverge as time goes on.

Conclusion

In this chapter, we have discussed the use of computational models for exploring possible mechanisms of change in cognitive development, focussing in particular on connectionist modelling. We have seen how these neurally inspired models of cognition can learn complex cognitive abilities when exposed to a training environment. Importantly, in theoretical terms, we have shown how these systems are not blank sheets, but contain constraints that shape the content that they can learn. We also explored how differences in the initial constraints under which a network develops can provide hypotheses about the causes of behavioural deficits in developmental disorders such as developmental dyslexia, autism, Specific Language Impairment, and Williams syndrome. This theoretical approach to explaining developmental deficits was contrasted with an earlier approach attempting to conceive of such deficits as equivalent to cases of adult brain damage. We finish the chapter with two final points.

First, it is important to clarify an idea that may have become obscured in the discussion of the many models outlined in this chapter. In any modelling endeavour, the aim is not solely to produce a model that can simulate human behaviour, whether typical or atypical. The aim is to derive an explanation of the target behaviour. Explanations based on connectionist models tend to be in terms of particular learning systems being exposed to cognitive domains with particular (statistical) structures, as we saw in the case of the balance scale and conservation models. However they must be clearly specified and empirically testable. It is not sufficient merely to point to a working model or list a set of connection strengths! Models allow the viability and coherence of theoretical ideas to be tested and, in so doing, drive theory forward. They also provide the opportunity to unify disparate empirical phenomena in a unified (and implemented) explanatory framework. But since models can be run in novel conditions, they also play an essential role in generating new testable empirical predictions, as we saw in the case of the models of seriation and of infants' expectancies concerning occluded objects.

Second, two major challenges remain for connectionist models of development in the future. The first challenge we alluded to earlier, in the need for more detailed accounts of the emergence of modules, and the need for simulations that address behaviour arising from cross-module interaction. The second challenge relates to training environments. The construction of a developmental model requires not just modelling the relevant cognitive system but also modelling the environment to which the system is exposed. The way in which connectionist models of development have focused attention on the exact nature of the environment has been one of their great strengths. However, the majority of current network models are passive recipients of their environment. In contrast, children are agents who, to a greater or lesser extent, control and seek out their learning environment. The next step for modelling will be to address the active nature of the child in development. This is of course a very complex issue. 'Dynamical' network systems, whose current performance affects their exposure to future training experiences, are liable to lapse into unstable, fluctuating states of temporary knowledge. It is probable that multi-component learning systems will be required to protect the system from this danger and create a relatively stable trajectory of development (see e.g. McClelland, McNaughton, & O'Reilly, 1995, for initial work on this idea).

Much work remains to be done in the modelling of development. However, connectionist networks provide an exceptionally useful tool for studying cognitive change – a tool that, as a servant of empirical investigation, can finally allow us to gain significant purchase on the role of the constraints of the mind and constraints of the environment in jointly driving cognitive development.

References

Baillargeon, R. (1993). The object concept revisited: New directions in the investigation of infants' physical knowledge. In C. E. Granrud (Ed.), *Visual perception and cognition in infancy* (pp. 265-315). LEA, London, UK.

Bellugi, U., Wang, P., & Jernigan, T. L. (1994). Williams syndrome: an unusual neuropsychological profile. In S. Broman and J. Grafman (Eds.), *Atypical cognitive deficits in developmental disorders: Implications for brain function* (pp 23-56). Hillsdale, NJ: Erlbaum.

Bishop, D. V. M. (1997). Cognitive neuropsychology and developmental disorders: Uncomfortable bedfellows. *Quarterly Journal of Experimental Psychology, 50A,* 899-923.

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

Bray, N. W., Reilly, K. D., Villa, M. F., & Grupe, L. A. (1997). Neural network models and mechanisms of strategy development. *Developmental Review*, *17*, 525-566.

Brown, G. D. A. (1997). Connectionism, phonology, reading, and regularity in developmental dyslexia. *Brain and Language*, *59*, 207-235.

Buckingham, D. & Shultz, T. R. (1994). A connectionist model of the development of velocity, time, and distance concepts. *Proceedings of the sixteenth annual Conference of the Cognitive Science Society* (pp. 72-77). Hillsdale, NJ: Erlbaum.

Bullinaria, J. A. (1997). Modelling reading, spelling, and past tense learning with artificial neural networks. *Brain and Language*, *59*, 236-266.

Christiansen, M. H., Allen, J. & Seidenberg, M. S. (1998). Learning to segment speech using multiple cues: a connectionist model. *Language and Cognitive Processes*, *13*, 221-268.

Christiansen, M. H. & Chater, N. (in press). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*.

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

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

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

Daugherty, K. & Seidenberg, M. S. (1992). Rules or connections? The past tense revisited. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 259-264). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*, 71-99.

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

Fahlman, S. E. & Lebiere, C. (1990). The cascade-correlation learning architecture. In D. S. Touretzky (Ed.), *Advances in neural information processing systems 2* (pp. 524-532). Los Altos, CA: Morgan Kaufmann.

Fodor, J. A. (1980). Fixation of belief and concept acquisition. In M. Piattelli-Palmarini (Ed.), *Language and learning: The debate between Chomsky and Piaget* (pp. 143-149). Cambridge, MA: Harvard Press.

Fodor, J. A. & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. In S. Pinker & J. Mehler (Eds.), *Connections and Symbols* (pp. 3-71). (*Cognition* Special Issue). Cambridge, MA: MIT Press / Bradford Books.

Forrester, N. & Plunkett, K. (1994). Learning the Arabic plural: the case for minority default mappings in connectionist networks. In A. Ramand &, K. Eiselt (Eds.), *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 319-323). Erlbaum, Mahwah, NJ.

Frawley, W. (in press). Control and cross-domain mental computation: Evidence from language breakdown. *Computational Intelligence*.
Gasser, M. & Smith, L. B. (1998). Learning nouns and adjectives: A connectionist account. *Language and Cognitive Processes, 13,* 269-306.
Gustafsson, L. (1997). Inadequate cortical feature maps: A neural circuit theory of autism. *Biological-Psychiatry, 42*, 1138-1147.
Hahn, U. & Nakisa, R. C. (in press). German inflection: Single route or dual-route? *Cognitive Psychology*.
Happé, F. G. E. (1994). *Autism: an introduction to psychological theory*. UCL Press.
Harm, M. W. & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review, 106*, 491-528.
Harris, N. G., Bellugi, U., Bates, E., Jones, W., & Rossen, M. (1997). Contrasting profiles of language development in children with Williams and Down syndromes. *Developmental Neuropsychology, 13*, 345-370.
Hoeffner, J. (1992). Are rules a thing of the past? The acquisition of verbal morphology by an attractor network. In *Proceedings of the 14th Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
Hoeffner, J. H. and 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 University Press.
Howlin, P., Davies, M. & Udwin, O. (1998). Cognitive functioning in adults with Williams syndrome. *Journal of Child Psychology and Psychiatry, 39*, 183-189.
Inhelder, B. & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence*. Basic Books: New York.
Jacobs, R. A. (1999). Computational studies of the development of functionally specialized neural modules. *Trends in Cognitive Science*, *3*, 31-38.
Jacobs, R. A., Jordan, M. I., & Barto, A. G. (1991). Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks. *Cognitive Science*, *15*, 219-250.
Johnson, M. H. (1999). Cortical plasticity in normal and abnormal cognitive development: Evidence and working hypotheses. *Development and Psychopathology, 11*, 419-437.
Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science.* MIT Press, Cambridge (Mass).
Karmiloff-Smith, A. (1997) Crucial differences between developmental cognitive neuroscience and adult neuropsychology. *Developmental Neuropsychology*, *13*, 513-524.
Karmiloff-Smith, A. (1998). Development itself is the key to understanding developmental disorders. *Trends in Cognitive Sciences*, *2*, 389-398.
Karmiloff-Smith, A., Grant, J., Berthoud, I., Davies, M., Howlin, P., & Udwin, O. (1997). Language and Williams syndrome: How intact is “intact”? *Child Development*, *2*, 246-262.
Karmiloff-Smith, A., Klima, E., Bellugi, U., Grant, J., & Baron-Cohen, S. (1995). Is there a social module? Language, face processing, and theory of mind in individuals with Williams syndrome. *Journal of Cognitive Neuroscience, 7*, 196-208.
Karmiloff-Smith, A., Tyler, L. K., Voice, K., Sims, K., Udwin, O., Howlins, P., & Davies, M. (1998). Linguistic dissociations in Williams syndrome: evaluating receptive syntax in on-line and off-line tasks. *Neuropsychologia*, *36*, 343-351.
Karmiloff-Smith, A. & Thomas, M. S. C. (in press). Developmental disorders. In M. A. Arbib (Ed.) *The Handbook of brain theory and neural networks*, 2nd Edition. Cambridge, MA: MIT Press.
Leonard, L. B. (1998). *Children with Specific Language Impairment*. Cambridge, MA: MIT Press.

Leslie, A., Xu, F., Tremoulet, P., & Scholl, B. (1998). Indexing and the object concept: Developing "what" and "where" systems. *Trends in Cognitive Sciences*, *2*, 10-18.
MacWhinney, B. & Leinbach, J. (1991). Implementations are not conceptualizations: revising the verb learning model. *Cognition*, *40*, 121-157.
Majerus, S., Palmisano, I., Van der Linden, M., Barisnikov, K., & Poncelet, M. (2001). An investigation of phonological processing in Williams syndrome. *Journal of the International*, *7(2)*, 153.
Marcus, G. (1998a). Can connectionism save constructivism? *Cognition*, *66*, 153-182.
Marcus, G. (1998b). Rethinking eliminative connectionism. *Cognitive Psychology*, *37*, 243-282.
Mareschal, D. (1991). Cascade-correlation and the Genetron: Possible implementations of Equilibration. *Technical Report 91-10-17*. McGill Cognitive Science Center, McGill University, Montreal, Canada.
Mareschal, D. (2000). Infant object knowledge: Current trends and controversies. *Trends in Cognitive Science*, *4*, 408-416.
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. R. (1999). Development of children's seriation: A connectionist approach. *Connection Science*, *11*, 153-188.
Mareschal, D. & Thomas, M. S. C. (in press). Self-organization in normal and abnormal cognitive development. In A. F. Kalverboer and A. Gramsbergen (Eds.), *Brain and Behavior in Human Development. A source book*, Dordecht: Kluwer Academic Publishers.
McClelland, J. L. (1989). Parallel distributed processing: implications for cognition and development. In R. G. M. Morris (Ed.) *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 9-45). Oxford University Press, Oxford, UK.
McClelland, J. L. & Jenkins, E. (1991). Nature, nurture, and connections: Implications of connectionist models for cognitive development. In K. van Lehn (Ed.), *Architectures for intelligence* (pp. 41-73) LEA, Hillsdale, NJ.
McClelland, J. L., McNaughton, B. L., & O'Reilly, R. G. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419-457.
McLeod, P., Plunkett, K. & Rolls, E. T. (1998) *Introduction to connectionist modelling of cognitive processes*. Oxford University Press, Oxford, UK.
Mervis, C., & Bertrand, J. (1997). Developmental relations between cognition and language: Evidence from Williams Syndrome. In L. B. Adamson & M. A. Romski (Eds.), *Research on communication and language disorders: Contributions to theories of language development* (pp. 75-106). New York: Brookes.
Milner, A. D. & Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford University Press, Oxford, UK.
Mitchell, T. M. (1997). *Machine learning*. McGraw-Hill: New York.
Morris, C. A., Demsey, S. A., Leonard, C. O., Dilts, C., & Blackburn, B. L. (1988). The natural history of Williams syndrome: physical characteristics. *Journal of Pediatrics*, *113*, 318-326.
Munakata, Y., McClelland, J. L., Johnson, M. N., & Siegler, R. S. (1997). Rethinking infant knowledge: Towards an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, *104*, 686-713.
Nakisa, R. C. and Plunkett, K. (1998). Innately guided learning by a neural network: the case of featural representation of speech language. *Language and Cognitive Processes*, *13*, 105-128.
Neville, H. J., Mills, D. L., & Bellugi, U. (1994). Effects of altered auditory sensitivity and age of language acquisition on the development of language–relevant neural systems: preliminary studies of Williams syndrome. In S. Broman and J. Grafman (Eds.), *Atypical cognitive deficits in developmental disorders: Implications for brain function* (pp. 67-83). Erlbaum.

O'Loughlin. C. & Thagard, P. (2000). Autism and coherence: A computational model. *Mind & Language*, *15*, 375-392.
O'Reilly, R. C. (1998). Six principles for biologically based computational models of cortical cognition. *Trends in Cognitive Sciences, 2,* 455-462.
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.
Pennington, B. F. & Smith, S. D. (1997). Genetic analysis of dyslexia and other complex behavioral phenotypes. *Current Opinion in Pediatrics, 9,* 636-641.
Piaget, J. (1954). *The Construction of Reality in the Child*. Basic Books, New York.
Piaget, J. (1965). *The child's concept of number*. Norton Library, New York.
Pinker, S. (1994). *The Language Instinct: The New Science of Language and Mind*. Penguin, New York.
Pinker, S. (1999). *Words and rules*. London: Weidenfeld & Nicolson
Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56-115.
Plunkett, K. (1998). Language acquisition and connectionism. *Language and Cognitive Processes*, *13*, 97-104.
Plunkett, K. & Juola, P. (1999). A connectionist model of English past tense and plural morphology. *Cognitive Science*, *23*, 463-490.
Plunkett, K. & Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition, 38,* 1-60.
Plunkett, K. & Marchman, V. (1996). Learning from a connectionist model of the acquisition of the English past tense. *Cognition*, *61*, 299-308.
Plunkett, K. & Nakisa, R. C. (1997). A connectionist model of the Arabic plural system. *Language and Cognitive Processes*, *12*, 807-836.
Plunkett, K. & Sinha, C. (1992). Connectionism and developmental theory. *British Journal of Developmental Psychology*, *10*, 209-254.
Plunkett, K., Sinha, C., Moller, M. F., & Strandsby, O. (1992). Vocabulary growth in children and a connectionist net. *Connection Science*, *4(3-4)*, 293-312.
Quartz, S. R. (1993). Neural networks, nativism, and the plausibility of constructivism. *Cognition*, *48*, 223-242.
Raijmakers, M. E., J., van Koten, S., & Molenaar, P. C. M. (1996). On the validity of simulating stagewise development by means of PDP networks: Application of catastrophe analysis and an experimental test of rule-like network performance. *Cognitive Science*, *20*, 101-136.
Reggia, J. A., Ruppin, E., & Berndt, R. S. (1996). *Neural modeling of brain and cognitive disorders*. London: World Scientific.
Rossen, M., Klima, E. S., Bellugi, U., Bihrle, A., & Jones, W. (1996). Interaction between language and cognition: Evidence from Williams syndrome. In J. H. Beitchman, N. J. Cohen, M. M. Konstantareas & R. Tannock, Language learning and behaviour (pp. 367-92). New York: Cambridge University Press.
Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 1 Foundations*. (pp. 318-362). Cambridge, MA: MIT Press.
Rumelhart, D. E. & McClelland, J. L. (1986). *Parallel Distributed Processing, Vol. 1*. MIT Press, Cambridge (Mass).
Rumelhart, D. E. & McClelland, J. L. (1986). On learning the past tense of English verbs. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 2: Psychological and*

Biological Models. (J. L. McClelland, D. E. Rumelhart, & the PDP Research Group), MIT Press, Cambridge (Mass), pp. 216-271.

Schafer, G. & Mareschal, D. (in press). Qualitative shifts in behavior without qualitative shifts in processing: the case of speech sound discrimination and word learning in infancy. *Infancy, 2*.

Schyns, P. G. (1991). A modular neural network of concept acquisition. *Cognitive Science*, *15*, 461-508.

Seidenberg, M. S. & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 452-477.

Shultz, T. R. (1998). A computational analysis of conservation. *Developmental Science*, *1*, 103-126.

Shultz, T. R., Buckingham, D., & Oshima-Takane, Y. (1994). A connectionist model of the learning of personal pronouns in English. In S. J. Hanson, T. Petsche, M. Kearns & R. L. Rivest (Eds.), *Computational learning theory and natural learning systems: Intersection between theory and experiment* (Vol. 2, pp. 347-362). Cambridge, MA: MIT Press.

Shultz, T. R., Mareschal, D., & Schmidt, W. C. (1994). Modeling cognitive development on balance scale phenomena. *Machine Learning, 16*, 57-86.

Siegler, R. S. (1981). Developmental sequences within and between concepts. *Monographs of the Society for Research in Child Development*, *46*, Serial No. 189.

Simonoff, E., Bolton, P., and Rutter, M. (1998). Genetic perspectives on mental retardation. In J. A. Burack, R. M. Hodapp, & E. Zigler (Eds.), *Handbook of mental retardation and development* (pp. 41-79). Cambridge University Press, Cambridge, UK.

Sirois, S. & Shultz, T. R. (1998). Neural network modelling of developmental effects in discrimination shifts. *Journal of Experimental Child Psychology*, *71*, 235-274.

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

Temple, C. (1997). *Developmental cognitive neuropsychology*. Psychology Press.

Temple, C. M., Almazan, M., & Sherwood, S. (in press). Lexical skills in Williams syndrome: a cognitive neuropsychological analysis. *Journal of Neurolinguistics*.

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

Thomas, M. S. C., Grant, J., Barham, Z., Gsödl, M., Laing, E., Lakusta, L., Tyler, L. K., Grice, S., Paterson, S. & Karmiloff-Smith, A. (2001). Past tense formation in Williams syndrome. *Language and Cognitive Processes*, *16(2/3)*, 144-176.

Thomas, M. S. C. & Karmiloff-Smith, A. (in press). Connectionist models of cognitive development, atypical development and individual differences. In R. J. Sternberg, J. Lautrey, & T. Lubart (Eds.), *Models of Intelligence for the Next Millennium*. American Psychological Association.

Thomas, M. S. C. & Karmiloff-Smith, A. (2001a). *Modelling language acquisition in atypical phenotypes*. Manuscript under revision.

Thomas, M. S. C. & Karmiloff-Smith, A. (2001b). *The implications of neural network models for causal explanations of behavioural impairments in cases of atypical development*. Manuscript in preparation.

Thomas, M. S. C. & Karmiloff-Smith, A. (2001c). *Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling*. Manuscript submitted for publication.

Udwin, O. & Yule, W. (1991). A cognitive and behavioural phenotype in Williams syndrome. *Journal of Clinical and Experimental Neuropsychology, 13*, 232-244.

Ullman, M. T., & Gopnik, M. (1999). Inflectional morphology in a family with inherited specific language impairment. *Applied Psycholinguistics*, *20*, 51-117.

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

Van der Lely, H. K. J. & Ullman, M. T. (2001). Past tense morphology in specifically language impaired and normally developing children. *Language and Cognitive Processes*, *16(2/3)*, 178-217.
Wilcox, T. (1999). Object individuation: Infants' use of shape size, pattern, and color. *Cognition*, *72*, 125-166.
Zorzi, M., Houghton, G., & Butterworth, B., (1998a). The development of spelling-sound relationships in a model of phonological reading. *Language and Cognitive Processes*, *13,* 337-372.
Zorzi, M., Houghton, G., & Butterworth, B. (1998b). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1131-1161.

Acknowledgements

Many thanks to Denis Mareschal for discussions of the ideas in this chapter, and Yonata Levy for helpful comments on an earlier draft. This work was funded by MRC Project Grant no. G9809880 and MRC Programme Grant no. G9715642 to Annette Karmiloff-Smith.