

Running head: CONNECTIONIST MODELS OF DEVELOPMENTAL DISORDERS

Connectionist Models of Developmental Disorders: A Critical Appraisal

Christopher M. Conway<sup>1</sup>, Michelle R. Ellefson<sup>2</sup>, Rick Dale<sup>1</sup>,  
and Morten H. Christiansen<sup>1</sup>

<sup>1</sup>Department of Psychology  
Cornell University

<sup>2</sup>Department of Psychology  
University of Warwick

### Abstract

This article critically reviews connectionist models of developmental cognitive disorders that are implemented by alterations to architecture, training dynamics, or input/output representations. Models are assessed in terms of four criteria: data contact, task veridicality, input representativeness, and deficit implementation. Data contact concerns the degree at which relevant empirical data is captured. Task veridicality assesses the extent to which a model's task corresponds to that of humans. Input representativeness refers to the degree to which the model's input matches the input available to a person. Finally, deficit implementation assesses the appropriateness of the network manipulations used to simulate a developmental deficit. We conclude that although some shortcomings exist, connectionist models show much potential for elucidating the underlying mechanisms of developmental cognitive disorders.

**Keywords:** developmental disorders; cognitive development; connectionism; computational models

## Connectionist Models of Developmental Disorders: A Critical Appraisal

### Introduction

One of the most important goals in cognitive science is to understand the development of cognition in relation to the changing organism and its environment. Although studying normal development may be the most direct approach towards accomplishing this goal, also it is common for researchers to rely on data from atypical developmental. Traditionally, the study of developmental disorders consists of examining dissociations in the cognitive profile of disordered versus normal populations (Temple, 1997). For instance, if a specific type of brain damage results in the same cognitive defect that is found in a developmental disorder, then one might conclude that the developmental disorder is caused by impairment to that particular brain region and/or cognitive module. However, a more recent, emerging perspective views developmental disorders in the context of trajectories and complex interactions between organism and environment (Elman et al., 1996; Thomas & Karmiloff-Smith, 2002). An example of this is found in Paterson, Brown, Gsödl, Johnson, and Karmiloff-Smith (1999) who report data regarding patients with Williams syndrome that highlights the importance of tracing cognitive deficits carefully throughout development. A related approach, focused on in this paper, is the use of computational models. Computational models provide an explicit instantiation of a theory that can then be tested and used to evaluate the theory's validity. The most common computational frameworks available for such simulations are symbolic and connectionist models. In regards to cognitive development, connectionist

models have been used more frequently than symbolic systems. This may be because connectionist networks provide a natural means for implementing mental/neural change, whereas symbolic models are less well suited for such applications (Bates & Elman, 1993).

Because connectionist models of cognition have grown rapidly in number and diversity, it can be difficult to evaluate their relative strengths and weaknesses. In this paper, we adopt the method taken by Christiansen and Chater (2001-a), who used three criteria to assess connectionist psycholinguistic models. We adapt these to evaluate connectionist models of developmental disorders, and invoke an additional criterion:

- *Data contact* refers to how well a model provides a fit with the relevant psychological data. We further distinguish between *primary* data contact, which refers to contact with data from specific experiments, and *secondary* contact, which refers to contact with general trends from a population. A model that exhibits primary data contact, for example, directly simulates dependent measures from psychological experiments, such as reaction time or performance accuracy in specific experimental conditions. Models that meet secondary data contact recreate general trends observed across experimental, observational, or anecdotal data. For example, a connectionist model may simulate poor memory performance in children with intellectual disabilities *generally*, without seeking to generate quantitative results from any specific study.
- *Task veridicality* assesses the extent to which a model's task is realistic relative to that which faces people. We further differentiate *training* task veridicality—the degree to

which a network's training task reflects what is faced by the population being simulated—from *test* task veridicality—the degree to which a network's test task reflects test conditions in real populations. Because of this distinction, a model that makes secondary data contact only can be assessed appropriately for its training task veridicality, whereas a model that makes primary data contact can be assessed for both training and test task veridicality. An example of good training task veridicality might be a model that acquires verb past tense morphology in a way that makes sense from the standpoint of how the language learner acquires past tense. An example of good test task veridicality might be a model that must produce the correct past tense form, given the word stem, if this is the task that faces the human participant as well.

- *Input representativeness* refers to the degree to which the information given to a model matches the input available to a person. For example, using natural language corpora to train language acquisition models is preferable to using artificial “toy” languages. Of course, no representational scheme can match up with reality perfectly. However, any scheme may be evaluated according to the extent to which it can be realistically attributed to some cognitive process in human subjects.
- *Deficit implementation* assesses how well the manipulations used to simulate a particular deficit are motivated by theoretical and/or empirical research. Karmiloff-Smith and Thomas (2003) outline several ways in which developmental disorders can be represented in connectionist networks, including modifications to the architecture, the training set, and the learning algorithm. An example of good deficit

implementation might be modifying a network's learning parameter to simulate delayed or stunted learning in children with intellectual disabilities. In this case, modifications to a parameter (learning rate) bears directly on the nature of the deficit being explored (learning speed or delay).

In this paper, we use these four criteria to evaluate a wide range of connectionist models of cognitive developmental disorders. It is important to note, however, that applying the criteria should take into account a model's relevant claims or goals. For example, we shall see below that oversimplification of network tasks or representations may sometimes serve as a useful shortcut for exploring development disorders. Often, such models are offered as *existence proofs* of the interaction of relevant variables, and how they may underlie a developmental disorder. For this reason, our four criteria offer a *flexible* guideline for evaluating connectionist models, and should be used in the context of a model's purported contribution to understanding what is being simulated.

The models in this review make contributions toward understanding cognitive developmental disorders, which we define simply as a cognitive deficit that is caused partly by genetic factors. Acquired disorders, such as may result from post-natal head trauma, peri-natal oxygen deprivation, or pre-natal alcohol exposure, are not included in this review. Furthermore, this paper focuses exclusively on developmental disorders that occur relatively early in life. Alzheimer's and Parkinson's diseases, as well as other developmental disorders that occur chiefly in post-adolescence, are excluded.

### Selecting a Deficit Implementation

Of crucial importance in modeling developmental disorders is choosing an appropriate method of implementing the underlying deficit. This is not always a straightforward enterprise, and it can be difficult to evaluate how well a model adequately achieves this. As an example, consider varying the size of a neural network's hidden layer as a way to implement a deficit where the hypothesized cause is reduced computational processing abilities. Despite being a well-motivated deficit implementation in terms of computational processing, this variation is not necessarily well motivated in terms of neural systems *per se* (i.e., the number of neurons, their interconnections, etc.). This illustrates the somewhat subjective nature of our criteria. When is a particular implementation considered to be adequately well-motivated? Complicating matters, connectionist models must balance realistic parameter manipulation against what can appear as a data-fitting practice (Roberts & Pashler, 2000). Though deficit implementation cannot be perfectly veridical in any model, a principled variation of hidden-unit layers to implement hypothesized computational processing deficits is at least a good place to start. Though this manipulation does not coincide with a direct neural realism, it should nevertheless *not* be considered data fitting.

To provide further illumination toward understanding the best way to implement a developmental disorder, we consider a recent neural network model (Oliver, Johnson, Karmiloff-Smith, & Pennington, 2000). Oliver et al. (2000) examined the emergence of normal and abnormal developmental trajectories by creating a two-layer, self-organizing neural network that resulted in clustered or topographic representations in response to input. To explore the effects that modifications of certain initial parameters may have on

development, Oliver et al. ran a series of simulations that varied several variables relating to the network's architecture, the training dynamics, and the input. Changes to the architecture included varying the excitatory node firing threshold and altering the distributions of node connections. Training dynamics were altered by modifying the connection decay rate. Finally, the input was modified by adjusting the amount of overlap that occurred between stimuli on the input layer. Modifications to these variables often resulted in large deficits in the emerging cortical representations relative to the baseline model. In fact, sometimes only a seemingly small change to these variables resulted in such deficits. For example, high or low threshold manipulations, or high decay rates, resulted in significantly less topographical organization.

Though not aimed directly at a specific disorder, these simulations provide a framework for considering different ways to implement developmental disorders. We adopt a similar deficit implementation framework and divide the ten models of this review into three categories based on implementation of the disorder: modifications to the model architecture, changes to the training dynamics, or variations to the input/output representations. With this review, we address two important points. First, we hope to demonstrate the motivation and benefits of using each type of deficit implementation in terms of the unique set of relations each has to neurophysiological or behavioral data. Second, we hope to uncover the strengths and weaknesses in current connectionist models of abnormal development with the hope that such an assessment will lead to new avenues of progress in the field. The remainder of this paper provides a critical appraisal of current connectionist models of developmental disorders, beginning with simulations relying on changes to the model's architecture.

## Architecture Manipulations

The first group of models implemented various disorders by modifying the networks' architectures. Included here are a model of autism (Cohen, 1998) and two models of developmental dyslexia: one of spelling (Brown & Loosemore, 1995) and one of surface dyslexia (Harm & Seidenberg, 1999). All three of these models implement the developmental disorders by varying the number of hidden units. Such an architectural modification can potentially be a useful way to model developmental disorders. For example, as we shall see, hypotheses concerning excessive neuronal cell growth or pruning can be approached abstractly through manipulating a model's architecture. Also, hypotheses concerning computational resources can be implemented by this technique.

### Cohen (1998): Varying hidden unit quantity to simulate abnormal cell density in the autistic brain

Cohen (1994, 1998) used a three-layer feedforward network with back-propagation (Rumelhart, Hinton, & Williams, 1986) to mimic learning and generalization deficits common in autism. An earlier model (Cohen, 1994) used input consisting of actual children's behavior patterns as reported in interviews with the children's caretakers. Based on this information, the model's task was to separate the cases into either autistic or mentally retarded. Unfortunately, this model does not address learning and behavior in autistic or mentally retarded children. Cohen's more recent model (1998) used similar theoretical assumptions but incorporated training and input representations to more

adequately investigate autism. Therefore, we will discuss the recent model exclusively below.

Autism occurs in roughly 2 to 5 children out of 10,000 (American Psychiatric Association, 1994), and results in persistent problems with language, cognition and social interaction (Cohen, 1998). Many neurological studies have been conducted to look for some biological etiology behind autism (see, for example, Bauman & Kemper, 1985, 1986; Kemper, 1988; Kemper & Bauman, 1993; Raymond, Bauman, & Kemper, 1989; Ritvo et al., 1986). One interpretation of these findings (Bauman, 1991) has been that autistic individuals have an increased cell density (“too many” neurons) or a decreased cell density (“too few” neurons) in certain areas of the brain. Some researchers have hypothesized that these brain abnormalities are caused by neuronal migration deficits (Piven et al., 1990) and/or abnormal neuronal growth (Bauman, 1991; Courchesne, Press, & Yeung-Courchesne, 1993). Based on these findings, Cohen simulated autism by varying the number of hidden units in the network. Cohen hypothesized that a model with too few or too many hidden units would qualitatively mimic the trend found in patients with autism, displaying difficulties with learning and generalization.

The model was used in service of a “fuzzy” (Kosko, 1993), nonlinear classification problem: vowel recognition in speech. This classification problem is regarded as fuzzy because the boundaries between vowels are subtle, depend on context, and vary across speakers. Autistic patients generally have difficulty mastering these types of problems. For example, they often display difficulties with semantic comprehension (Kanner, 1943) or taking other points of view (Frith & Baron-Cohen, 1987). In the present model, Cohen used frequency formant data from five different

vowels pronounced by individuals varying in gender and age (Nowlan, 1990). The frequency formant data contains fuzzy boundaries because there is substantial overlap among the different vowels; in addition, the same vowel is pronounced differently by different speakers. The network's task was to learn to classify each vowel using the formant input data. A set of vowels that the networks had not seen before was used for subsequent testing. The overall results indicated that networks with too many hidden units could learn the classification task but had difficulties generalizing to the new formant data. Conversely, networks with too few hidden units performed poorly at both classification and generalization. An optimal number of hidden units allowed for the best performance during both the learning and test phase.

These results mimic a general trend in autism, most notably that of the inability to generalize patterns to novel instances (Lovaas, Schreibman, Koegel, & Rehm, 1971; Schreibman & Lovaas, 1973). For example, Cohen noted that some individuals with autism have difficulty understanding verbal requests if the request is not worded in a specific, familiar way. Cohen discussed how these and other simulation results might account for other characteristics of children with autism, such as: difficulty in learning complex higher-order concepts, resistance to change, inattention to socially relevant stimuli, stimulus over-selectivity, and highly successful pattern discrimination.

Brown & Loosemore (1995): Varying hidden unit quantity to simulate reduced computational resources in dyslexia

In another model making use of modifications to network architecture, Brown and Loosemore (1995) simulated spelling development in normal children and children with

dyslexia. Developmental dyslexia is a disorder characterized by reading levels substantially below average despite normal intellectual functioning and environmental conditions. Difficulties range from spelling to recognizing real words to sounding out non-words with regular orthography to phonology pronunciation. Several theories have been put forth to explain dyslexia (e.g., Bradley & Bryant 1983; Bryant & Bradley 1985; Frith 1985; Stanovich 1988; Torgesen 1999), but most center around the learner's knowledge of and ability to master language's phonological structure.

Brown and Loosemore observed that those with dyslexia are poor at non-word reading and spelling relative to control subjects at the same reading level (Rack, Snowling, & Olson, 1992). However, as they argued elsewhere (Brown & Watson, 1991), children with dyslexia do not show reduced regularity effects. Brown and Loosemore hypothesized that these traits—impaired non-word spelling with little deficits to regularity effects—could be characterized by reducing the amount of computational resources available to a network during training. Their reasoning was based on Seidenberg and McClelland (1989), who simulated the differences between skilled and poor young readers by a reduction in the number of hidden units.

Similar to the Cohen (1998) model, a three-layer feedforward network was trained using standard backpropagation learning (Rumelhart et al., 1986). The task of the network was to use the phonological form of a monosyllabic word and output the corresponding spelling. This model used a “triple based representation”, or wickelfeatures, for both orthography and phonology. The phonological and orthographic representations of words were transformed into feature triples that were then given to the model as input (Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989). For

example, the word *mint* has four wickelfeatures; *\_mi*, *min*, *int*, and *nt\_* (where “\_” signifies a word boundary). The input and output layers each contained 50 units. Brown and Loosemore used 35 hidden units to train normal performance, but reduced this number to 20 and 15 hidden units to implement mild and severe dyslexia, respectively.

The input consisted of words that conformed to either a regular or irregular spelling pattern. Nineteen tokens of each category were used. To represent spelling regularity, each regular word was accompanied by an average of four words that were similar in orthography and phonology (e.g., the regular word *can* was accompanied by *pan* and *tan*). Each irregular word was accompanied by an average of four words that were equivalent in phonology but dissimilar in orthography (e.g., the irregular word *rain* was accompanied by *mane* and *lane*). A total of 223 words were presented as input to each of the networks. The network was evaluated in terms of its error in producing the correct orthography given a word’s phonology. Models with smaller hidden unit layers, as an obvious consequence, learned more slowly and never achieved the same level of accuracy as the normal model.

Finally, the networks were tested on the spelling of regular and irregular novel non-words. To do this, the authors matched all three models on their performance for regular words. The three models had the same performance on regular word spelling at different points in training: the normal model after 130 epochs, the mildly dyslexic model after 390 epochs, and the severely dyslexic model after 1580 epochs. With each of the three models matched on regular words, the authors compared performance on non-words and on irregular words. They found that all of the models performed equally on irregular words but non-word performance was severely worsened depending on the dyslexic

severity. The performance on the non-words conformed to general experimental data regarding dyslexic performance (e.g., Rack, Snowling, & Olson, 1992). The authors conducted a separate experiment to examine the reading performance of 24 children with dyslexia matched on reading level with 24 control children. As predicted by their model, they found no difference between dyslexics and non-dyslexics on irregular (or regular) word spelling. Thus, the authors believe that developmental dyslexia might arise from a reduction in the amount of computational resources available during learning.

Harm & Seidenberg (1999): Varying clean-up and hidden unit quantities to simulate phonological impairments in dyslexia

In the final model we present in this section, Harm and Seidenberg (1999) manipulated network architectures, as well as training dynamics, to simulate two subtypes of developmental dyslexia, surface and phonological. Surface dyslexia is characterized by greater deficits in reading exception words (e.g., *pint*) than non-words (e.g., *tilk*). On the other hand, phonological dyslexia is characterized by greater deficits in reading non-words rather than exception or irregular words. Harm and Seidenberg created an attractor network that learned the structure of phonemic segments by using phonetic features in slots corresponding to phonemes. The network represented phonological knowledge in terms of attractors in state space. This allowed the model to fill in missing features, but it was limited in how it represented phonological mappings. The network was trained to read in two stages. In the first stage units representing the phonetic features of words were trained to “auto-associate” word patterns by utilizing a

“cleanup” unit layer, which corrected any errors. Next the network was trained using back-propagation to map orthography to phonology.

Learning to read occurred in three different conditions: (a) trained attractor condition, (b) untrained attractor condition, (c) simple feedforward network. During the simulation of normal reading development, the connections between phonetic feature units and the cleanup apparatus were eliminated leaving only connections from orthography to hidden units and from hidden units to phonological units. Generally, the phonological attractor architecture facilitated learning the orthography-phonology mapping task, but phonological knowledge did not have to be in place prior to reading acquisition because it could be acquired rapidly along the way.

After demonstrating the applicability of their general model to reading acquisition, Harm and Seidenberg manipulated the architecture and training dynamics of the model to simulate phonological dyslexia. Phonological dyslexia was simulated by impairing the phonological information available to the model before training. Mild phonological impairment was implemented by decaying the weights in the phonological attractors (training dynamics manipulation). More severe impairment was achieved by removal of the phonological clean up units (architecture manipulation). More specifically, this consisted of continued weight decay and the severing of a random selection of half of the connections running between phonological units.

The models' phonological impairments were assessed by performance on a pattern completion task. The results of network performance indicated that increased amounts of damage resulted in decreased ability to complete the task. The degree of impairment was indicative of the pattern of human performance in relation to the exception (e.g. *choir*)

and non-words (e.g., *domb*). The models suffering mild impairment did not have diminished performance levels on the exception words, while the non-words were only slightly affected. This result is very similar to data from individuals with pure phonological dyslexia. Additional impairment yielded a "mixed" pattern of errors. In these models, performance on non-words was impaired. Although these results replicated the patterns of results in human studies, the error rates in the exception words were lower than actual error rates of reading in children (e.g., Manis et al., 1996).

An additional group of models simulated the effects of phonological impairments on speech perception. The models were trained using a standard discrimination task with similar stimuli as an empirical study using children with dyslexia (Werker & Tees, 1987). Simulations with mild impairments elicited a significant negative impact on speech perception. More severe impairments resulted in larger deficits in task performance. Although the mild impairments match empirical data, the severe impairment results resemble specific language impairment (SLI) more so than a dyslexic pattern of results.

Next, the authors simulated surface dyslexia by another architecture manipulation, varying the number of hidden units. With fewer hidden units, the networks had a significant performance reduction on the reading of exception words. With more severe deficits (even fewer hidden units) the model's ability to generalize also became impaired, further decreasing the network's performance on regular verbs, simulating a decreased capacity to encode orthography to phonology mappings. This model simulated what the authors referred to as a "mixed surface dyslexia" because the primary deficits were found in performance of exception words with some deficits also occurring with non-words.

In summary, the architecture and training dynamics manipulations implemented by Harm and Seidenberg appear to match the behavioral evidence. Damaging the models' capacity to adequately represent phonology for words resulted in deficits in the rate of learning, the level of performance, and the developmental pathway. By altering the number of hidden units in the models, they were able to simulate both mild and severe dyslexia.

### Evaluation

This section reviewed three sets of connectionist simulations dealing with autism (Cohen, 1998), spelling impairments in developmental dyslexia (Brown & Loosemore, 1995), and reading impairments in phonological and surface dyslexia (Harm & Seidenberg, 1999). Cohen's simulations demonstrated a possible relation between number of neurons and learning and generalization deficits in the autistic brain. Specifically, with too few hidden units, the model was poor at learning, while with too many hidden units, the model was poor at generalization. By the criteria outlined in the introduction, Cohen's model matches up quite well. The secondary data contact by this model is rather good, since it matches broadly established behavioral evidence. Also, by using frequency formant information, the simulations approach rather realistic kinds of information relevant to speech processing. Though the task itself of classifying vowels is rather simple, it is similar to the type of task children may face. Deficit implementation at face value is also good, because, if varied neuronal density is the cause for autism as has been suggested, then modifying the number of hidden units is a well-motivated implementation of this disorder.

The Brown and Loosemore model successfully approaches primary data contact: the model accounts for non-word and irregular spelling performance in dyslexics. The training and test veridicality exhibited by the model is adequate, since mapping from orthography to sound is an essential component of the common portrait of reading skill; and the test task for the network fits quite well with that of the children in the human experiment. In fact, the same stimuli were used for both the model and the experiment. However, the model seems to have drawbacks in input representativeness and deficit implementation. First, input to the network consisted of a very small vocabulary. Also, there are well-known problems with wickelfeatures, which have been criticized because they result in a loss of information, low generalizability, and non-linguistic patterns (e.g., Pinker & Prince, 1988). Thus, this model could be improved by incorporating input representations that better match what is assumed to underlie true phonological knowledge. Finally, the deficit implementation is not motivated by well-established neurological or psychological evidence, making the reduction of hidden units (i.e., computational resources) less clear as an appropriately veridical means to approach dyslexia. As we mentioned in the introduction, however, this idealization may serve the simulations as a useful principled demonstration of a more abstract cognitive deficit. Additionally, as Zorzi, Houghton, and Butterworth (1998) point out, there is neuroimaging evidence that some forms of dyslexia are related to an inability to allocate neural resources toward reading tasks (Hynd & Semrud-Clikeman, 1989). If so, then the reduction of hidden units may be a justified deficit implementation for dyslexia.

Finally, the Harm and Seidenberg simulations as a whole are quite good in terms of these criteria. The models make good primary data contact: the results of the

simulations were comparable to the behavioral evidence of phonological and surface dyslexia (Manis et al., 1996). Input representativeness is decent, because the simulations make use of distinctive features commonly used to describe underlying phonological knowledge. Both training and test task veridicality are also good. In particular, the authors attempted to train the model in two stages to better represent how learning to read depends on prior learning with phonological information. And finally, the deficit implementation is well-motivated because the authors used theoretical constructs of both phonological and surface dyslexia, based on empirical data, to create the parameters of their models.

Both the Cohen (1998) and Harm and Seidenberg (1999) models meet our four criteria quite well. Brown and Loosemore's (1995) simulations could be improved by using more realistic input representativeness as well as rethinking the method of deficit implementation. All three models in this section used modifications to the hidden unit layer to implement the various disorders. Of the three, the Cohen model's deficit implementation could be considered the most well-motivated because it was based on neural evidence of the autistic brain.

### Changes to Training Dynamics

This second section reviews three connectionist models that implement developmental disorders by making changes to the networks' training dynamics. Included are models of intellectual disabilities (Bray, Reilly, & Grupe, 1997), schizophrenia (Hoffman & McGlashan, 1997), and autism (O'Laughlin & Thagard, 2000). By training

dynamics, we refer to parameters that may be modified in a neural network during learning. For example, as we shall describe below, simply observing neural network performance during differing stages of training may offer clues about the nature of a developmental disorder. Also, as another example, changing connection weights during training, such as by pruning or adding noise, may permit exploration of different patterns of development.

Bray et al. (1997): Freezing weights to simulate a delayed or frozen developmental trajectory in children with intellectual disabilities

Bray, Reilly, and Grupe (1997) used neural network models to explore the mechanisms involved in the development of external memory strategies by children with and without intellectual disabilities (mental retardation). External memory strategies refer to strategies that are non-verbal or gesture-related in nature (e.g., pointing, touching, or manipulating items). For instance, placing one's briefcase next to the door is an external reminder to remember to bring the briefcase to work. Such strategies are useful in that they reduce the demands on the cognitive system. Behavioral evidence indicates that external memory strategies are ubiquitous throughout the lifespan, for 2- and 3-year-olds (DeLoache, Cassidy, & Brown, 1985; Wellman, Ritter, & Flavell, 1975), older children (Kruetzer, Leonard, & Flavell, 1975), and adults (Intons-Peterson & Fournier, 1986). Some philosophers have even argued that such external strategies should be considered part of our cognitive system proper (Clark & Chalmers, 1998).

Bray, Saarnio, Borges, and Hawk (1994) used an "instruction following task" to test for differences in the development of external memory strategies in 7- and 11-year-

old children without intellectual disabilities and 11-year-old children with intellectual disabilities. In this experiment, the children were presented with 12 small movable objects (e.g., a shoe, a pencil, a key) and 6 fixed toy targets (a toy chair and toy refrigerator) and heard sequences of instructions informing them what to do with the objects (e.g., “put the shoe on the chair” and “put the comb in front of the refrigerator”). A baseline condition tested to see whether children would spontaneously employ external memory strategies to help them remember the set of instructions. In an experimental condition, the children were either prompted to use external memory strategies or instructed specifically to use a particular strategy. Bray et al. (1994) recorded the type of strategy that the children displayed as well as their accuracy in performing the task. The children’s strategies were divided into two general categories: (a) object-oriented strategies, involving pointing at or manipulating an object without orienting toward the target; (b) target-oriented strategies that involved orienting an object toward the specific target. The results showed that the type of external strategy used by the children affected their overall performance, namely that target-oriented strategies helped performance while object-oriented strategies appeared to hinder performance. In the baseline condition, the 7-year-old children without intellectual disabilities and the 11-year-old children with intellectual disabilities used the object-oriented strategy more than the older children without intellectual disabilities. The reason for this may be because as children develop, they begin to rely less on external strategies and more on internal, verbal strategies. In the experimental condition, the prompt and instructions resulted in all groups of children increasing their use of the more-efficient target-oriented strategies.

Bray et al. (1997) used a model named the Generalized Components/Attention Bias Model (GC/AB) to simulate this empirical data. The GC/AB is a hybrid model, composed of a connectionist implementation of a modular architecture (Hrycej, 1992), using local, rather than distributed, representations. Several components comprise the GC/AB, each module representing a particular aspect of strategy behavior: sequencer and associative memory modules represent, store, and recall a sequence of events; a strategy module connects to the associative memory nodes, influencing recall depending on the strategy type; an attention/bias module influences strategy selection; an accuracy-feedback module represents an "external teacher", keeping track of whether recall was correct; and a tactics module simulates the ability of children to encode more information and more complex strategies with increased development. The instructions given to the children were coded as input sequences to the model, in terms of the appropriate object, target, and relation (e.g., "on" or "in front of"). Using a Hebbian learning algorithm, the network's task was to choose one of the three memory strategies to encode the instructions. The network produces output that reveals the emergence of strategy development (Bray et al., 1997): over the course of the network's "life", it relies more extensively on the more complex and efficient strategy type.

Previous simulations had shown that the GC/AB model captures general trends in the development of these strategies in normally developing children (Anumolu, Bray, & Reilly, 1997). Of more importance to our present purpose, the model was also used to simulate the development of memory strategy selection in 11-year-old children with intellectual disabilities. In order to simulate the performance of these impaired children, Bray et al. established a "criterion run" for the model, based on normal 11-year old

performance. Twelve different networks were trained for 32 epochs, each network corresponding to one child, and each epoch representing a single instruction sentence. The criterion run of the model simulated the strategy use for 11-year-olds without intellectual disabilities in the control condition. Then, the performance of the model on the criterion run was analyzed to find the epoch that provided the best fit to data of the children with intellectual disabilities. The connection weights for this trial were frozen and used to simulate the pattern of connectivity for the 11-year-old children with intellectual disabilities in the control condition. This same procedure was also used to create a network that provided a best fit for the 7-year-old children without intellectual disabilities. The 11<sup>th</sup> epoch provided the best fit for these 7-year-olds while the 14<sup>th</sup> epoch was used to simulate the children with intellectual disabilities. Finally, these three networks simulated the empirical data of the three groups of children in the experimental condition. The simulation results provided a good fit with the behavioral data for control children and children with intellectual disabilities. Bray, Villa, Reilly, Grupe, and Anumolu (2000) suggested that these results show that the impaired children are on the same developmental trajectory as the older children, only at a less advanced stage, which they believe resonates with current developmental theories of intellectual disabilities<sup>1</sup>.

Hoffman & McGlashan (1997): Pruning hidden/context unit weights to simulate atypical brain connectivity in schizophrenia

We next consider computational work devoted to schizophrenia. Although schizophrenia typically is regarded as a psychotic and not a developmental disorder

---

<sup>1</sup> For a different perspective, see Thomas & Karmiloff-Smith (2002).

(American Psychiatric Association, 1994), certain evidence discussed below suggests that atypical brain development may be an underlying cause and therefore warrants inclusion in this review. Four lines of inquiry have been conducted using computational models to explore this idea further. Cohen and Servan-Schreiber (1992) used a connectionist model to explore schizophrenic deficits in terms of impairments to parameters representing the effects of dopamine in the brain. Hoffman and Dobscha (1989) and Hoffman and McGlashan (1993) created an attractor network (Hopfield, 1982) to examine the relation between cortical pruning and various schizophrenic symptoms. Horn and Ruppin (1995) and Ruppin, Reggia, and Horn (1996) also used an attractor network to explore how changes in synaptic strength may lead to schizophrenic delusions and hallucinations. The last series of investigations, performed by Hoffman and McGlashan (1997) used a simple recurrent network (SRN; Elman, 1990) to explore the origin of auditory hallucinations in schizophrenia. We describe this study more fully below.

It is well established that synaptic density in the normal cortex peaks at childhood and then declines before reaching adult levels (e.g., Huttenlocher, 1979); this process is known as synaptic pruning. Hoffman and McGlashan (1997) hypothesized a connection between such declines in cortical connectivity and schizophrenia (Feinberg, 1982/1983; Hoffman & Dobscha, 1989). Specifically, they created a neural network model of speech perception in order to explore the hypothesis that schizophrenia is a result of irregularities in synaptic pruning during development. The model focuses on auditory hallucinations, which are common in schizophrenia (Sartorius, Shapiro, & Jablonsky, 1974).

The SRN received input vectors that were used to represent abstract phonetic patterns for a vocabulary of 30 words (14 nouns, 11 verbs, 4 adjectives, and the auxiliary

contraction, *won't*). The output neurons coded words in terms of semantic and syntactic features. For example, output nodes represented the word “cop” by coding for features such as “noun”, “animate”, “human”, and others. Thus, the network’s task was to associate input phonology with a constellation of syntactic and semantic features.

Training comprised 60 repetitions of 256 unique grammatical sentences using a variant of back-propagation learning (Miikkulainen, 1993). During the test phase, the network was exposed to a set of 23 novel sentences using the same vocabulary. Each test sentence was separated by a pause of five null inputs, accomplished by setting each input node to zero. Hallucinations were said to occur if the network produced output features for a word during one of these pauses, when no phonetic input was present.

Because prior studies had shown that schizophrenics display working memory impairments (Gold, Carpenter, Randolph, Goldberg & Weinberger, 1997; Park & Holzman, 1992; Weinberger, Berman, Suddath & Torrey, 1992), Hoffman and McGlashan targeted the hidden unit and context layers of their model by applying the concept of “neural Darwinism” (Edelman, 1987). That is, if a connection weight linking the context and hidden layers fell below a specified threshold, it was set permanently to zero. Weight pruning was varied from between 0% and 95%.

The results revealed that the standard (0% pruning) network correctly detected 67% of test words. But as the hidden layer connections underwent pruning, word detection increased up to a maximum of 91% (occurring at 64% pruned connections). With additional pruning (between 80% and 95%), word detection rapidly decreased and hallucinations were produced. Hallucinations consisted of the single word *won't*, occurring immediately after nouns that it had followed during training. Thus, as the

researchers remark, the hallucinations were not random, but conformed to linguistic constraints specific to the training materials.

Hoffman and McGlashan provided some additional support for their theory that cortical connectivity reductions could lead to schizophrenic symptoms in humans. They observed that the network performed best at a 64% level. If one considers this level of pruning to represent normal adult levels of connectivity, it is interesting to note that an additional 20% reduction in the network corresponds to the onset of auditory hallucinations. This estimate matches neurophysiological findings of the density of frontal cortex in schizophrenic versus normal brains (Garey et al., 1995; Selemon, Rajkowska, & Goldman-Rakic, 1995). Although this post-hoc analysis walks the line of prediction and data fitting, the model arguably provides insight into the nature of schizophrenic hallucinations in relation to synaptic pruning.

O’Laughlin & Thagard (2000): Modifying excitatory/inhibitory strengths to simulate impaired information integration in autism

We turn now to another model of autism, this one exploring the role of processing multiple levels of information. The cognitive deficits found in autism, according to Frith’s (1989) weak coherence theory, at least in part are due to an inability of children with autism to successfully integrate various levels of information. Consequently, children with autism must process pieces of information one-at-a time and have difficulty creating coherent representations of the world around them, resulting in a reduction of their cognitive and social skills (Frith & Happé, 1994). O’Laughlin and Thagard (2000) developed a connectionist model to further investigate how an inability to process

multiple levels of information might influence the performance of children with autism on a false belief and a homograph task.

In a typical False Belief Task (Wimmer & Perner, 1983), a child watches as an object, perhaps a marble, is unexpectedly moved while a person is absent from the room. After the object is moved, the child is asked where the departed person might look for the marble upon returning to the room. In order to correctly solve the task, the child must integrate information about the current location of the marble and the absent person's belief about the object's location. More specifically, children must inhibit the most available choice (the current location of the object) in order to select the correct, but less available choice, of the object's original location. Previous research has indicated that children under about 4 years of age and children with autism have difficulty correctly solving this task (Baron-Cohen, Leslie, and Frith, 1985). The authors used Thagard's ECHO (Thagard, 1992) – a connectionist model that encodes explanatory and contradictory connections between propositions – to simulate the abilities of children with autism to incorporate the multiple levels of information required to solve the false belief task. The model, as well as children faced with this task, must integrate the appropriate information in order to reach the correct solution.

The authors tested three types of models on the false belief task. The various models tested how manipulating the strength of connections between nodes influenced how the network arrived at a solution. One possibility is that if the excitatory weights are much larger than the inhibitory weights, then the network may bounce back and forth and never arrive at a solution. Because the inhibitory links between nodes would be small

compared to larger excitatory links, the network would not be able to settle on a solution because it would continually be considering illogical responses.

Alternatively, a network could arrive at a solution prematurely. This response might be elicited by large inhibitory weights. In this case the model ignores the plausible response by inhibiting its activation. If the plausible response is not activated then the network will settle on a premature solution because it is stuck in a local minimum. In a connectionist network, weights are modified in relation to the error of the system. Based on the amount of error, appropriate changes are made to continually reduce the error. If the weight changes are not large enough, then the network may incorrectly arrive at a solution in the local minimum, rather than in the actual error minimum, called the global minimum. The authors proposed that failure to cohere in this type of task might be a result of the network settling prematurely in a local minimum.

The first type of simulation was based on empirical findings with young children. Initially, young children will pay more attention to the most available choice: the current location of the marble. In order to simulate this type of response the authors created a connectionist model with excitatory connection strengths for the most available answer and inhibitory connection strengths for the correct, but less available answer: the original location of the object. The authors suggested that this simulation is analogous to the performance of children younger than 4 years of age. The model, similar to a young child, was biased to use the most recent information. In this case, the model was most recently exposed to the change in location of the marble. Based on the coherence theory, the authors assumed that older children more efficiently cohere relevant information. To simulate the performance of older children, the strength of the connections between the

nodes representing the knowledge of the absentee person and the chosen response of looking at the first location were strengthened. Once the model could simulate the performance of children over 4 years of age, the authors again modified the weights to attempt to simulate the inability of older children with autism to cohere relevant information. In this simulation the inhibitory connection strengths were increased, thus making it more difficult for the model to select the correct choice. Increasing the inhibition strengths resulted in an imbalance of the inhibitory and excitatory connections in the network. As a result, the network prematurely settled into a local minimum. Because access to the correct solution was not easily available, the network was unable to cohere the multiple levels of information and arrived prematurely at the incorrect solution.

A second type of task, the homograph task, was simulated to assess the influence of context to solve a problem necessitating coherence. In the English language homographs are words with the same spelling that require different pronunciations and have a different semantic definition, depending on the context surrounding the word (e.g., a *tear* in the eye vs. a *tear* in a shirt). Often the preferred pronunciation of a homograph is biased by the relative frequency of possible pronunciations. In this example, English speakers pronounce *tear* more often as a *tear in the eye* than as a *tear in a dress*. Children with autism have difficulty selecting the correct pronunciation of homographs (e.g., Frith & Snowling, 1983; Happé, 1997). To arrive at the correct pronunciation of *tear in a dress*, the model must incorporate the appropriate context and inhibit the more frequent pronunciation. If the network cannot integrate the relevant information, it will be unable to reach the correct solution. To simulate performance in the homograph task,

the authors used a connectionist model called IMP (Kunda & Thagard, 1996). The IMP model is an associative network of positive and negative connections between various elements that represent a typical context for *tear* in a dress. Similar to the false belief simulations, the IMP model was given excitatory connection weights for the pronunciation that is most frequent and inhibitory connection strengths for the less frequent pronunciation. Under these initial training conditions, the model was able to successfully inhibit the frequency effect, integrate context, and arrive at a correct solution. As in the false belief task simulation, increased inhibition resulted in a disregard of context and premature arrival at the incorrect solution.

In summary, increasing the amount of inhibitory strength compared to excitatory strength resulted in an inability of the model to correctly solve both the false belief and the homograph task. The results of the model support Frith's theory that an inability to cohere multiple sources of information negatively affects performance on the false belief and homograph tasks. In addition, the results suggest that increases in inhibition, relative to excitation, might be responsible for the impaired abilities of children with autism to cohere information.

### Evaluation

This section illustrated three ways in which aspects of training dynamics can be varied to explore the nature of developmental disorders. The Bray et al. (1994) model achieves a fairly good level of data contact in their model of intellectual disability. In particular, the primary data contact is quite good. However, the deficit implementation for reaching this level of data contact is questionable: it does not seem likely that

intellectual disability is adequately simulated by freezing the weights of the normal model corresponding to the epoch which provides a best-fit. In fact, if each epoch corresponds to an instruction sequence, then freezing the model's weights on the 11th epoch translates into halting the child's task after the 11th sentence. Additionally, this implementation appears to suggest that intellectual disability follows a normal developmental path until a certain point, when an abrupt change occurs either in the form of a halt in learning or input presentation. This does not seem to be the case for individuals with intellectual disabilities. Perhaps a more appropriate method of deficit implementation would be to manipulate the learning rate parameter in the impaired model, which would result in a gradual, ongoing deficit, rather than the unrealistic case of an abrupt and sudden onset.

In general, Hoffman and McGlashan's (1997) network did not substantially reflect the pattern of auditory hallucinations that schizophrenics experience. The authors admitted that actual hallucinations often consist of whole phrases or sentences, not merely one word, and do not occur solely in response to external speech, as the network showed. Also, the model predicts that schizophrenics are poor at word detection; it is not clear that this prediction is substantiated. For these reasons, the model's secondary data contact appears inadequate. In addition, the input corpus itself provides a highly idealized representation of actual language, and similarly, the task of mapping phonetic patterns onto syntactic and semantic features is a considerable abstraction. Despite these limitations, implementing the deficit by weight pruning is well-motivated by evidence of the relations between synaptic density, working memory, and schizophrenia.

As for O’Laughlin and Thagard (2000), there were data on the performance of children with autism in a false belief and homograph task, but the authors did not compare the results of their simulations to a specific set of results. However, because the models predicted decisions that were similar to the responses of children with autism for these same two tasks, they provide good secondary data contact. The paper could have been considerably strengthened if the authors had included a quantitative comparison between the performance of the models and children with autism. Another disadvantage is that the model was hand-coded, rather than relying on a learning algorithm. As such, it is questionable as to what the model can tell us about developmental trajectories, and might be more accurately regarded as an exercise in data fitting. There were some good aspects of the model, however. For instance, the model was provided the same type of information and was required to make the same decisions children faced during the same task, therefore leading to good input representativeness and task veridicality. Finally, the deficit implementation is also good. The motivation behind modification of the connection strengths between inhibition and excitation was based on a relevant theory concerning the possible reason for a decline of cognitive abilities in children with autism.

### Changes to Input/Output Representations

The final group of models we consider modify network representations to implement developmental disorders. In the following models, this technique was used to model specific language impairment (SLI; Hoeffner & McClelland, 1993), dyslexia (Brown, 1997), and Williams Syndrome (WS; Thomas & Karmiloff-Smith, 2003-a). As

we shall see, modifying input representations, such as by modifying the number of input/output units, is a convenient and well-controlled technique for altering the nature of sensory input to a neural network. This fact makes the technique suitable for exploring disorders potentially related to perceptual information.

Hoeffner & McClelland (1993): Reducing input strength to simulate a perceptual processing deficit in SLI

Hoeffner and McClelland (1993) explored morphology acquisition within the context of developmental dysphasia, also known as SLI. SLI is the diagnostic category for children who display abnormal language development (e.g., deficits in phonology, morphology, or syntax) but are otherwise seemingly normal in all other respects. Some propose that SLI is a genetic, language-specific impairment (Gopnik, 1997; Gopnik & Crago, 1991; Pinker, 1994) whereas others theorize that SLI is caused by a perceptual processing deficit (Krauss et al., 1996; Leonard, 1988; Leonard, 1993; Tallal, 1990; Tallal et al., 1996; for recent summaries of SLI research, see Bishop, 1997; Joanisse & Seidenberg 1998; Leonard, 1997). Adopting the perceptual processing deficit hypothesis, Hoeffner and McClelland presented a three-layer, attractor neural network account of SLI in the domain of verb morphology.

The model, originally developed to simulate processing in normals (Hoeffner, 1992), had three layers: phonological, hidden, and semantic. The model's task was to learn mappings between semantic and phonological representations. The corpus used for training, based on Marcus et al. (1992), comprised 1925 monosyllabic regular and irregular verb forms divided into five categories: unmarked form (*run*), past tense (*ran*),

third-person singular (*runs*), progressive (*running*), and past participle (*run*). The input phonology was represented by eight slots: three prevocalic, one vocalic, three postvocalic, and one syllabic suffix, with each containing ten units corresponding to standard phonological features. In the standard model, an input strength (i.e., degree of activation received from a set of “perceptual” units) of 30 was used for all of the phoneme/position combinations except for word-final stops and fricatives, which had an input strength of 10. Word-final phonemes started out with less weight than the other phonemes in both models corresponding to what occurs in natural speech: word-final phonemes are weakly stressed, dropped or assimilated into the next word. The networks were trained using the Contrastive Hebbian Learning algorithm (Peterson & Anderson, 1987), consisting of two phases, one in which the network was trained on both the phonological and semantic inputs, and one in which only the semantic input was used.

The SLI deficit was modeled in a second network that was trained using the same number of presentations as the standard model. To simulate SLI according to the perceptual processing deficit, all phoneme strengths were reduced by one-third, including the already weakly weighted word-final stops and fricatives. Relative to the standard model, the impaired model captured several key deficits of SLI, including: slower overall learning with more errors; more “no marking errors” (e.g., the unmarked form, *walk*, is produced instead of *walked* or *walks*); and more deficits for regular past-tense learning than irregular past-tense learning. The model showed a lack of uniform impairment for words with morphemic vs. nonmorphemic endings. For example, the model correctly produced *need* 100% of the time while *died* was produced correctly only 10% of the time. This is because perceptual salience interacts with competition between verb form

categories for a particular item. *Need* does not have to compete with *nee*, but *died* does have to compete with *die*. This explanation answers previous critiques of the perceptual processing account, namely that all words must be equally affected (Gopnik & Crago, 1991).

Brown (1997): Using a distributed input representation to simulate phonological impairments in dyslexia

In another model using input modifications, Brown (1997) built on the earlier work of Brown and Loosemore (1995), discussed previously, to further examine the “phonological deficit” hypothesis. This hypothesis claims that developmental dyslexia results from an impairment to an orthography-to-phonology processing mechanism (see Snowling, 1987, for a review). The standard view of the phonological deficit hypothesis leads to two predictions. The first is that dyslexics will be impaired at non-word reading (e.g., *slint*). The second is that they will show reduced spelling-to-sound regularity effects in real words. That is, pronunciation for an “irregular” word (e.g., *pint*) will be difficult for normal children (e.g., because of the tendency to use the same pronunciation as that used for *mint* and *hint*). However, because of their phonological processing deficit, children with dyslexia will not share this same effect. The first prediction has been confirmed by many experimental studies (see Rack, Snowling, & Olson, 1992, for a review). However, numerous studies have failed to confirm the second prediction (e.g., Seidenberg, Bruck, Fornarolo, & Backman, 1985; Treiman & Hirsh-Pasek, 1985). Thus, Brown (1997) offered a connectionist simulation in an attempt to explore the

contradiction between the predictions of the phonological deficit hypothesis and the actual experimental data.

Brown used two different networks to simulate the development of reading in normal children and those with dyslexia. The standard model, based on Plaut, McClelland, Seidenberg, and Patterson (1994), used a localist representation scheme. Each input node corresponded to a particular letter and was activated when a word containing those letters was used as input (e.g., if *pint* was the input, nodes corresponding to *p*, *i*, *n*, and *t* were activated). Groups of units encoded the onset (initial consonants), vowel, and coda (final consonants) of each word. The second model, which was based on the models of Seidenberg and McClelland (1989) and Brown and Loosemore (1995), employed a distributed input representation that was used to simulate the phonological deficit hypothesis in developmental dyslexia. Previous computational work has demonstrated that such a system of representation lacks important properties for proper generalization. As in Brown and Loosemore (1995), this model represented orthography and phonology with wickelfeatures (e.g., *pint* is represented by *\_pi*, *pin*, *int*, and *nt\_*). Both of the present models were three-layer, feed-forward networks, containing 35 hidden units, and trained via back propagation (Rumelhart et al., 1986).

For both models the vocabulary was similar to that used above in Brown and Loosemore (1995). The input consisted of 19 regular, consistent words (e.g., *pill*), 19 irregular/exception words (e.g., *pint*), and 189 other words that were selected so as to provide context for the regular and irregular words (e.g., *pill* was accompanied by *hill* and *mill*; *pint* was accompanied by *mint* and *lint*). The models learned to output a

phonological representation of the input. After training, the models were tested on their reading performance for regular words, irregular words, and non-words.

The simulation results revealed that the dyslexic model, as predicted by the phonological deficit hypothesis, was impaired at non-word reading. However, the simulations also showed that the dyslexic and standard models had equivalent orthography-to-phonology regularity effects. This second result counters the prediction of the phonological deficit hypothesis but agrees with much of the experimental literature, as mentioned above. Finally, Brown conducted an additional experiment using 10 children with severe dyslexia matched on reading level with 10 non-dyslexic controls. Participants read words that appeared on a computer screen as quickly as possible. Errors in pronunciation were recorded in terms of the percentage of regulars and irregulars. Children with dyslexia and controls had very similar performance--suggesting the same orthography-to-phonology regularity effects. The experimental results clearly supported the data gained from the simulation, but were not consistent with the standard phonological deficit hypothesis. Brown concluded that the connectionist approach provides the best interpretation of the phonological deficit hypothesis.

Thomas & Karmiloff-Smith (2003-a): Increasing input/output unit quantities to simulate an over-emphasis on phonology in Williams Syndrome

Thomas and Karmiloff-Smith (2003-a) explored the atypical development of past tense in patients with Williams Syndrome (WS). WS is a genetic neurodevelopmental disorder that is characterized by an uneven cognitive-linguistic profile (Howlin, Davies, & Udwin, 1998; Mervis, Morris, Bertrand, & Robinson, 1999): patients often have

average verbal abilities with poor visuospatial abilities. Thomas and Karmiloff-Smith used Plunkett and Marchman's (1993) model of past tense formation and explored five different manipulations embodying five different hypotheses regarding the nature of WS. The five hypothesis-based models were tested in relation to the three most critical features of past tense performance in children with WS: (a) high performance on regular verbs, (b) lower performance on irregular verbs, and (c) poor generalization.

Each of the models simulated by Thomas and Karmiloff-Smith were trained and tested using the same basic materials and procedure. The materials were the same as the language of Plunkett and Marchman (1993). This artificial language contained non-word verb stems that were created using three phonotactic templates of English (CVC, CCV, and VCC). Past tense could be formed regularly or irregularly. Regular past tense was formed by adding /d/ (e.g., *tame* → *tamed*), /t/ (e.g., *wrap* → *wrapped*), or /ed/ (e.g., *wait* → *waited*). There were three types of irregular past tense formation: arbitrary (e.g., *go* → *went*), no change (e.g., *hit* → *hit*), and vowel change (e.g., *ring* → *rang*). To test the ability of the model to generalize, novel verbs were presented. These novel verbs rhymed with existing regular verbs, rhymed with existing irregular verbs, or did not rhyme with any existing verbs.

By manipulating the amount of detail available at both the input and output level, Thomas and Karmiloff-Smith examined the effects of an overly detailed representation of phonology. According to this hypothesis, the language system in a child with WS places more importance on the phonological information and less importance on semantic information. Although much of the evidence supporting this hypothesis is circumstantial, there is some empirical evidence that children with WS are over-sensitive to auditory

stimuli (e.g. Finegan, Smith, Meschino, Vallance & Sitarenious, 1995; Mervis & Bertrand, 1997; Neville, Holcomb, & Mills, 1989; Neville, Mills & Bellugi, 1994). In order to test this hypothesis, the authors allowed five different levels of phonological detail to the models. This was implemented by increasing the number of input and output units allotted for each phoneme based on its context in a word, leading to greater discrimination of phonemes. For example, the /d/ in *read* becomes more distinct from the /d/ in *made* as the phonological information becomes more detailed. The model with the most detailed phonological representation simulated existing empirical data for the three most critical characteristics of WS. Large amounts of phonological detail produced a delay in the development of both regular and irregular past tense formations. Moreover, increases in phonological detail produced decreases in phonological similarity, which resulted in a decrease of the model's ability to generalize. Increasing the detail of the phonological representations inevitably increased the amount of information available to the network. In order to insure that the results of the simulations were not influenced by this increase in information, additional models controlled for redundancy by randomly introducing noise.

In subsequent models, the authors implemented other types of manipulations, including learning rates, lesions to network architectures, number of hidden units, simulation parameters, noise levels, number of internal layers, plasticity, and level of unit discriminability to simulate past tense learning in children with WS. None of these subsequent manipulations sufficiently simulated the critical features of WS. Some manipulations were slightly promising in that they were able to simulate two of the critical features (e.g., manipulating the learning rate). Some models, which held

promising theoretical bases, did not adequately simulate the empirical data (e.g., manipulating redundancy, lesioning the exception route, and modifying the amount of hidden units, layers, plasticity, and noise). However, Thomas and Karmiloff-Smith explored the effect of a semantic-phonological integration deficit by changing the lexical-semantic information in the model. This model, impaired in its ability to integrate phonology and semantics, also made contact with the WS data.

### Evaluation

In this final section, we considered ways in which input/output representations may be modified to explore developmental disorders. In view of our criteria, Hoeffner and McClelland (1993) make good contact with secondary data, although it is important to note that the model does not take into account other linguistic problems associated with SLI children (e.g., syntactic difficulties; Joanisse & Seidenberg, 1998). Task veridicality, input representativeness, and deficit implementation are also rather good. The semantics-to-phonology mapping task used for training seems to be a reasonable approximation of an important part of acquiring morphology. Additionally, reducing the strength of word final phonemes appears to be a valid way to represent what occurs in natural speech. Finally, based on the perceptual processing impairment hypothesis, reducing the phonological strength to all phonemes is a well-motivated deficit implementation.

Brown's (1997) simulations captured experimental data showing that children with dyslexia have deficits in reading non-words, but when matched on reading level, show an equal orthography-to-phonology regularity effect. These simulations also used a

reading task like that faced by the children. It is therefore successful in primary data contact and task veridicality. Deficit implementation is also good, the model being based on the phonological deficit hypothesis, which is one theory as to the cause of developmental dyslexia. The input representativeness, however, is a relative weakness due to the reasons discussed with the Brown and Loosemore (1995) model (e.g., small vocabulary sizes and the use of wickelfeatures).

Thomas and Karmiloff-Smith's simulations set out to examine the developmental disorder of WS and explore how adjustments in the models' specifications impact their ability to acquire past tense forms. Focusing specifically on the overly detailed phonological representation model, they achieve good primary data contact. This model simulated all of the critical empirical findings of WS. However, it is lacking in training task veridicality because it is unlikely that children learn past tense morphology by acquiring stem to past tense mappings (e.g., Akhtar and Tomasello, 1997, demonstrate that verb inflection acquisition is more complicated than this). We should point out, though, that this task does fit within common modeling practice (Bullinaria, 1997). The model's test task veridicality is good because the models were tested on tasks that closely simulate the experiments from which the empirical data was derived (Thomas et al., 2001). Specifically, both the networks and human participants were given a word stem and were required to produce the correct past-tense form. The authors were able to examine a number of theoretical constructs underlying the empirical data regarding children with WS and examined whether or not certain hypotheses adequately described the developmental nature of WS. Therefore, the models are rather good examples of deficit implementation.

## Discussion

As suggested in the introduction, there are a number of important theoretical and methodological details that must be taken into account when evaluating connectionist models of developmental disorders. We consider some of these below, with respect to each of our criteria. No model, of course, can perfectly satisfy all of them together or even individually. In fact, as mentioned in the introduction, it may be argued that flouting these potential virtues may provide idealizations that maintain tractability and thereby more clearly contribute to understanding variables underlying developmental disorders. However, scaling up is a broader problem that deserves attention. Thomas and Karmiloff-Smith (2003-a) have demonstrated that multiple deficit implementations may lead to the same disorders in connectionist models – suggesting that decisions regarding appropriate implementations are more complicated than usually assumed. For this reason, models that incorporate more realistic input representations, have greater task veridicality, and more closely approach primary and secondary data, could provide added insight into the appropriate deficit implementations. In the remainder of the Discussion, we summarize the overall assessment of the models reviewed in this paper.

*Data contact* was our first assessment criterion, which we used to evaluate how well the results of each model were supported by available psychological evidence. Primary data contact refers to data from actual human experiments whereas secondary data contact concerns general evidence from the target population not in the form of specific cognitive/behavioral experiments. It is clear that most of the models exhibit

good contact with data. However, only half of the models make primary data contact. Although we suggest that primary data contact should be the goal for computational models of cognition, it may not always be appropriate to do so. For example, Hoeffner and McClelland (1993) achieve excellent secondary data contact in their simulation of SLI, illustrating an important theoretical point about perceptual salience in a verb production task. In this case, primary data contact would not necessarily strengthen the model with respect to the authors' broader theoretical goal. However, developmental dyslexia has been approached through several models, many of which achieve primary data contact. This could be a sign that primary data contact may be necessary to offer thorough computational insights regarding the variables contributing to developmental disorders.

We note that there is, of course, more to capturing psychological data than merely fitting empirical results. A good model (or theory) also should make non-obvious predictions and must be sufficiently constrained to be falsifiable through empirical tests (see Roberts & Pashler, 2000, for further discussion). Certainly, in the case of a young field, often it is desirable first to model general trends of the target population and then only later attempt to simulate data from actual human experiments.

*Task veridicality* and *input representativeness* both varied widely among the models. The former was used to assess the extent to which each model simulated the relevant behavioral and cognitive tasks. The latter assesses the degree to which the input to a model corresponds to what would be available to a person. In general, it appears that models of developmental dyslexia rate consistently high for both task veridicality and

input representativeness. The other models have a wider range of performance with these two criteria.

Finally, the last criterion used in this paper was *deficit implementation*. This measure assesses the network modifications that were used to simulate each respective disorder. Here, we find that most of the models perform quite well. It may be that connectionists have focused mainly on the implementation of the deficit itself, at the expense of the task and input. We are also reminded of Thomas and Karmiloff-Smith (2003-b) who underscored the complexity of selecting an appropriate deficit implementation, potentially complicating the evaluation of theories of developmental disorders. However, as mentioned above, we suggest that scaling up in terms of our criteria may help constrain models sufficiently so that the most appropriate deficit implementation becomes apparent. This in turn will lead to a more complete understanding of developmental disorders in general.

Case in point, three of the models we reviewed targeted developmental dyslexia, each model using a different deficit implementation. Because all three manipulations appear to produce dyslexic-like behavior in the models, what kinds of conclusions can we make regarding the underlying causes of the disorder? First, one should consider that it is certainly possible that a particular impairment could be simulated by more than one model manipulation, and as such it should not come as a surprise to see multiple deficit implementations for the same disorder. In addition, if a particular hypothesized impairment is successfully simulated in multiple ways, this might give added support to the hypothesized impairment. In the present example, both Brown (1997) and Harm and Seidenberg (1999) implemented dyslexia by simulating a phonological impairment,

lending added support to a causal link between phonological impairment and dyslexia. Finally, we suggest that more weight should be given to the model that most satisfactorily meets the four criteria. In our judgment, because Harm and Seidenberg (1999) best incorporate realistic input, task, and contact with data, the manner in which they implemented the deficit is likely to provide the most insight into the actual disorder.

### Conclusion

This paper has provided a critical review of connectionist models of developmental disorders. We reviewed three groups of models on the basis of their method of deficit implementation, including changes to architecture, training dynamics, and input/output representations. Table 1 summarizes the models and how each simulated the disorder in question. We find that architecture (hidden unit layer) modifications are particularly well-suited for simulating deficits relating to atypical neuron densities and/or limitations to computational processing. Modifications to a model's training dynamics appear useful for exploring the effects of deficits that change dynamically throughout development. Finally, manipulations to input/output representations are appropriate for modeling perceptual processing deficits. Though modelers of developmental disorders ought to strive for appropriate deficit implementations, this review has emphasized the importance of providing additional constraints in the form of realistic data contact, input, and task veridicality. Scaling up a model in terms of these criteria may provide a more substantial framework with which to investigate the underlying factors related to any particular developmental disorder.

In closing, we note that there seem to be more models dealing with linguistic disorders than non-linguistic disorders (for additional discussion concerning modeling developmental language disorders, see Bullinaria, 1997; Manis, Seidenberg, Doi, McBride-Chang, & Petersen, 1996; Zorzi, Houghton, & Butterworth, 1998). In addition, those dealing with linguistic phenomena tend to better satisfy our criteria. Why would models of language deficits be more prevalent and better constructed than those dealing with non-language disorders? One possible answer is that there has been a close marriage between language processing and connectionist models. Indeed, “connectionist psycholinguistics” (Christiansen & Chater, 2001-b) has emerged as a prominent field in the past decade, and it has provided a firm foundation upon which models of abnormal development can be built. Regardless of the reason for this imbalance, we anticipate that future research will address the weaknesses in the models of both linguistic and non-linguistic disorders. Specifically, we suggest that more attention needs to be paid toward making good primary data contact, incorporating relevant tasks, and creating realistic input. We also agree with Fischer and Connell (2000), who argue that models of a developmental disorder should strive to capture key properties of a developmental pathway, taking into account the various components affecting development, rather than merely simulating a single point on the pathway. Current models of developmental disorders do not meet this standard adequately.

Despite these shortcomings, we find that the models reviewed indicate that connectionism may hold great promise for modeling developmental cognitive disorders. This promise is underscored by the fact that many researchers are beginning to view abnormal development in terms of connectionist principles. For example, Nation (1999)

proposes the use of connectionist modeling to develop a more advanced theoretical framework for hyperlexia, a condition that is often associated with autism. Snowling, Hulme, and Goulandris (1994) use a connectionist framework to interpret the developmental deficits found in the case study of a dyslexic boy. And Van Orden, Bosman, Goldinger, and Farrar (1997) outline an unimplemented recurrent network account that emphasizes the role of phonology representation in developmental dyslexia. These examples, and others, show that connectionism is growing in popularity as a conceptual framework for thinking about cognitive developmental disorders. In this context, it is probably not a coincidence that two of the models in this review that met our criteria the best (Harm & Seidenberg, 1999; Thomas & Karmiloff-Smith, 2003-a) are also two of the most recent ones. We believe that connectionist principles offer continuing promise for elucidating the mechanisms underlying developmental cognitive defects and hope that our modeling criteria may help realize this potential.

## References

- Akhtar, N. & Tomasello, M. (1997). Young children's productivity with word order and verb morphology. *Developmental Psychology*, *33*, 952-965.
- American Psychiatric Association (1994). *Diagnostic and statistical manual of mental disorders* (4<sup>th</sup> ed.). Washington, DC: Author.
- Anumolu, V., Bray, N.W., & Reilly, K.D. (1997). Neural network models of strategy development in children. *Neural Networks*, **10**(1), 7-24.
- Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a 'theory of mind'? *Cognition*, *21*, 37-46.
- Bates, E.A. & Elman, J.L. (1993). Connectionism and the study of change. In M.H. Johnson (Ed.), *Brain development and cognition*. Cambridge, MA: Basic Blackwell.
- Bauman, M. (1991). Microscopic neuroanatomic abnormalities in autism. *Pediatrics*, *87*, 791-796.
- Bauman, M., & Kemper T.L. (1985). Histoanatomic observations of the brain in early infantile autism. *Neurology*, *35*, 866-874.
- Bauman, M., & Kemper T.L. (1986). Developmental cerebellar abnormalities: A consistent finding in early autism. *Neurology*, *36*, 190.
- Bishop, D.V.M. (1997). *Uncommon understanding: Development and disorders of language comprehension in children*. Hove, East Sussex, U.K: Psychology Press.

- Bradley, L. & Bryant, P.E. (1983). Categorizing sound and learning to read: A causal connection. *Nature*, *301*, 419-421.
- Bray, N.W., Reilly, K.D., & Grupe, L.A. (1997). Neural network models and mechanisms of strategy development. *Developmental Review*, *17*, 525-566.
- Bray, N.W., Saarnio, D., Borges, J.M., & Hawk, L.W. (1994). Intellectual and developmental differences in external memory strategies. *American Journal on Mental Retardation*, *99*, 19-31.
- Bray, N.W., Villa, M., Reilly, K.D., Grupe, L.A., & Anumolu, V. (2000). *The Generalized Components/Attention Bias neural network model of strategy development in children with and without mental retardation* [On-line]. Available: <http://www.uab.edu/cogdev/sandiego.htm>.
- Brown, G. D. A. (1997). Connectionism, phonology, reading, and regularity in developmental dyslexia. *Brain and Language*. *59*, 207-235.
- Brown, G.D.A. & Loosemore, R. (1995). A computational approach to dyslexic reading and spelling. In C.K. Leong & R.M. Joshi (eds.), *Developmental and acquired dyslexia*, (pp. 195-219). Kluwer Academic Publishers. Dordrecht, The Netherlands.
- Brown, G.D.A. & Watson, F.L. (1991). Reading in developmental dyslexia: A connectionist approach. In M. Snowling & M. Thomas (Eds.), *Dyslexia: Integrating theory and practice* (pp. 165-182). Whurr Publishers: London.
- Bryant, P.E. & Bradley, L. (1985). *Children's reading problems*. Oxford: Blackwell.
- Bullinaria, J.A. (1997). Modeling reading, spelling, and past tense learning with artificial neural networks. *Brain and Language*, *59*, 236-266.

- Cazden, C.B. (1968). The acquisition of noun and verb inflections. *Child Development*, 39, 433-448.
- Christiansen, M.H. & Chater, N. (2001-a). Connectionist psycholinguistics: Capturing the empirical data. *Trends in Cognitive Sciences*, 5, 82-88.
- Christiansen, M.H. & Chater, N. (2001-b). Connectionist psycholinguistics in perspective. In M.H. Christiansen & N. Chater (Eds.), *Connectionist psycholinguistics*, (pp. 19-75). Westport, CT: Ablex.
- Clark, A. & Chalmers, D.J. (1998). The extended mind. *Analysis*, 58, 10-23.
- Cohen, I.L. (1994). An artificial neural network analogue of learning in autism. *Biological Psychiatry*, 36, 5-20.
- Cohen, I.L. (1998). Neural network analysis of learning in autism. In D.J. Stein & J. Ludvik (Eds.), *Neural networks and psychopathology*, (pp. 274-315). Cambridge: Cambridge University Press.
- Cohen, J.D. & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behaviour and biology in schizophrenia. *Psychological Review*, 99, 45-77.
- Courchesne, E., Press, G.A., Yeung-Courchesne, R. (1993). Parietal lobe abnormalities detected with MR in patients with infantile autism. *American Journal of Roentgenology*, 160, 387-393.
- DeLoache, J.S., Cassidy, D.J., & Brown, A.L. (1985). Precursors of mnemonic strategies in very young children's memory. *Child Development*, 56, 125-137.
- Edelman, G.M. (1987). *Neural Darwinism*. New York: Basic Books.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.

- Elman, J.L., Bates, E.A., Johnson, M.H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Feinberg, I. (1982/1983). Schizophrenia: Caused by a fault in programmed synaptic elimination curing adolescence? *Journal of Psychiatric Research*, *17*, 319-334.
- Finegan, J.-A., Smith, M. L., Meschino, W. S., Vallance, P. L., & Sitarenious, G. (1995). *Verbal memory in children with Williams syndrome*. Indianapolis: Society for Research in Child Development.
- Fischer, K.W. & Connell, M.W. (2000). Using neuroconstructivist tools to understand developmental pathways to disorders. *Developmental Science*, *3*, 24-26.
- Frith, U. (1985). The usefulness of the concept of unexpected reading failure: Comments on "Reading retardation revisited." *British Journal of Developmental Psychology*, *3*, 15-17.
- Frith, U. & Baron-Cohen, S. (1987). Perception in autistic children. In D.J. Cohen & A.M. Donnellan (Eds.), *Handbook of autism and pervasive developmental disorders* (pp. 85-102). New York: Wiley.
- Frith, U., & Happé, F. (1994). Autism: Beyond 'theory of mind'. *Cognition*, *50*, 115-132.
- Frith, U., & Snowling, M. (1983). Reading for meaning and reading for sound in autistic and dyslexic children. *British Journal of Developmental Psychology*, *1*, 329-342.
- Garey, L.J., Ong, W.Y., Patel, T.S., Kanani, M., Davis, A., Hornstein, C., & Bauer, M. (1995). Reduction in dendritic spine number on cortical pyramidal neurons in schizophrenia. *Abstracts of the Society for Neuroscience*, *21*, 237.

- Gold, J.M., Carpenter, C., Randolph, C., Goldberg, T.E., & Weinberger, D.R. (1997). Auditory working memory and Wisconsin Card Sorting Test performance in schizophrenia. *Archives of General Psychiatry*, *54*, 159-165.
- Gopnik, M. (1997). Language deficits and genetic factors. *Trends in Cognitive Sciences*, *1*, 5-9.
- Gopnik, M., & Crago, M. (1991). Familial aggregation of a developmental language disorder. *Cognition*, *39*, 1-50.
- Happé, F. G. E. (1997). Central coherence and theory of mind in autism: Reading homographs in context. *British Journal of Developmental Psychology*, *15*, 1-12.
- Harm, M.W., & Seidenberg, M.S. (1999). Phonology, reading acquisitions, and dyslexia: Insights from connectionist models. *Psychological Review*, *106*, 491-528.
- Hoeffner, J. (1992). Are rules a thing of the past? The acquisition of verbal morphology by an attractor network. In *Proceedings of the 14<sup>th</sup> annual meeting of the cognitive science society* (pp. 861-866). Hillsdale, NJ: Erlbaum.
- Hoeffner, J.H. and McClelland, J.L. (1993). Can a perceptual processing deficit explain the impairment of inflectional morphology in development dysphasia? A computational investigation. In E.V. Clark (Ed.), *The proceedings of the 25th annual child language research forum* (pp. 38-49). Stanford University, Center for Study of Language and Information.
- Hoffman, R.E. & Dobscha, S. (1989). Cortical pruning and the development of schizophrenia: A computer model. *Schizophrenia Bulletin*, *15*, 477-490.
- Hoffman, R.E. & McGlashan, T.H. (1993). Parallel distributed processing and the emergence of schizophrenic symptoms. *Schizophrenia Bulletin*, *19*, 119-140.

- Hoffman, R.E. & McGlashan, T.H. (1997). Synaptic elimination, neurodevelopment, and the mechanism of hallucinated “voices” in schizophrenia. *American Journal of Psychiatry*, *154*, 1683-1689.
- Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, *79*, 2554-2558.
- Horn, D. & Ruppin, E. (1995). Compensatory mechanisms in an attractor neural network model of schizophrenia. *Neural Computation*, *7*, 182-205.
- Howlin, P., Davies, M., & Udwin, O. (1998). Syndrome specific characteristics in Williams syndrome: To what extent do early behavior patterns persist into adult life? *Journal of Applied Research in Intellectual Disabilities*, *11*, 207-226.
- Hrycej, T. (1992). *Modular learning in neural networks: A modularized approach to neural network classification*. New York: Wiley.
- Huttenlocher, P.R. (1979). Synaptic density in the human frontal cortex – developmental changes and effects of aging. *Brain Research*, *163*, 195-205.
- Hynd, G. W. & Semrud-Clikeman, M. (1989). Dyslexia and brain morphology. *Psychological Bulletin*, *106*, 447-482.
- Intons-Peterson, M.J., & Fournier, J. (1986). External and internal memory aids: When and how often do we use them? *Journal of Experimental Psychology: General*, *115*, 267-280.
- Joanisse, M.E. & Seidenberg, M.S. (1998). Specific language impairment: A deficit in grammar or processing? *Trends in Cognitive Sciences*, *2*, 240-247.
- Kanner, L. (1943). Autistic disturbances of affective contact. *Nervous Child*, *2*, 217-250.

- Karmiloff-Smith, A. & Thomas, M.S.C. (2003). Developmental disorders. In M.A. Arbib and P.H. Arbib (Eds.), *Handbook of brain theory and neural networks* 2<sup>nd</sup> Edition. MIT Press.
- Kemper, T.L. & Bauman, M.L. (1993). The contribution of neuropathologic studies to the understanding of autism. *Behavioral Neurology*, *11*, 175-187.
- Kemper, T.L. (1988). Neuroanatomic studies of dyslexia and autism. In J.W. Swann & A. Messer (Eds.), *Disorders of the developing nervous systems: Changing views on their origins* (pp. 125-154). New York: Alan R. Liss.
- Kosko, B. (1993). *Fuzzy thinking: The new science of fuzzy logic*. New York: Hyperion.
- Krauss, N., McGee, T.J., Carrell, T.D., Zecker, S.G., Nicol, T.G., & Koch, D.B. (1996). Auditory neurophysiologic responses and discrimination deficits in children with learning problems. *Science*, *273*, 971-973.
- Kruezer, M.A., Leonard, C. & Flavell, J. H. (1975). An interview study of children's knowledge about memory. *Monographs of the Society for Research in Child Development*, *40*, (Whole No. 159).
- Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, *103*, 284-303.
- Leonard, L. (1993). The use of morphology by children with specific language impairment: Evidence from three languages. In R. Chapman (Ed.), *Processes in language acquisition and disorders*. St. Louis: Mosby-Yearbook.
- Leonard, L. (1997). *Children with specific language impairment*. Cambridge, MA: MIT Press.

- Leonard, L., Sabbadini, L., Volterra, V., & Leonard, J. (1988). Some influences on the grammar of English and Italian speaking children with specific language impairment. *Applied Psycholinguistics*, 9, 39-57.
- Lovaas, O.I., Schreibman, L., Koegel, R.L., & Rehm, R. (1971). Selective responding by autistic children to multiple sensory input. *Journal of Abnormal Psychology*, 77, 211-222.
- Manis, F., Seidenberg, M., Doi, L., McBride-Chang, C., & Paterson, A. (1996). On the basis of two subtypes of developmental dyslexia. *Cognition*, 58, 157-195.
- Marcus, G.F., Pinker, S., Ullman, M., Hollander, M., Rosen, T.J., & Xu, F. (1992). Overregularization in language acquisition. SRCO Monograph, Serial No. 228, Vol. 57, No. 4.
- 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 communications and language disorders: Contributions to theories of language development* (pp. 75-106). New York: Brookes.
- Mervis, C. B., Morris, C. A., Bertrand, J., & Robinson, B. F. (1999). Williams syndrome: Findings from an integrated program of research. In Helen Tager-Flusberg et. al. (Eds.) *Neurodevelopmental disorders* (pp. 65-110). Cambridge: MIT Press.
- Miikkulainen, R. (1993). *Subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory*. Cambridge, MA: MIT Press.
- Nation, K. (1999). Reading skills in hyperlexia: A developmental perspective. *Psychological Bulletin*, 125, 338-355.

- Neville, H. J., Holcomb, P. J., & Mills, D. M. (1989). Auditory, sensory and language processing in Williams syndrome: An ERP study. *Journal of Clinical and Experimental Neuropsychology*, *11*, 52.
- 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. 63-83). Mahwah, NJ: Erlbaum.
- Nowlan, S.J. (1990). *Competing experts: An experimental investigation of associative mixture models*. Technical Report CRG-TR-90-5, University of Toronto.
- 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*, 1-40.
- O'Loughlin, C., & Thagard, P. (2000). Autism and coherence: A computational model. *Mind & Language*, *15*(4), 375-392.
- Park, S. & Holzman, P.S. (1992). Schizophrenics show spatial working memory deficits. *Archives of General Psychiatry*, *49*, 975-982.
- Paterson, S.J., Brown, J.H., Gsödl, Johnson, M.H., & Karmiloff-Smith, A. (1999). Cognitive modularity and genetic disorders. *Science*, *286*, 2355-2358
- Peterson, C. & Anderson, J.R. (1987). A mean field theory learning algorithm for neural networks. *Complex Systems*, *1*, 995-1019.
- Pinker, S. (1991). Rules of language. *Science*, *253*, 530-535.

- Pinker, S. (1994). *The language instinct: How the mind creates language*. New York: William Morrow.
- Pinker, S. (1999). *Words and rules*. London: Weidenfeld & Nicolson
- Pinker, S. & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28, 73-193.
- Piven, J., Berthier, M.L., Starkstein, S.E., Nehme, E., Pearlson, G., Folstein, S. (1990). Magnetic resonance imaging evidence for a defect of cerebral cortical development in autism. *American Journal of Psychiatry*, 147, 734-739.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1994). *Understanding normal and impaired word reading: Computational principles in quasi-regular domains*. Carnegie Mellon University. [Technical report PDP.CNS.94.5]
- Plunkett, K. & Marchman, V. (1993). From rote learning to system building: acquiring verb morphology in children and connectionist nets. *Cognition*, 48, 21-69.
- Rack, J.P., Snowling, M.J., & Olson, R.K. (1992). The nonword reading deficit in developmental dyslexia: A review. *Reading Research Quarterly*, 27, 29-51.
- Raymond, G., Bauman, M., & Kemper T. (1989). The hippocampus in autism: Golgi analysis. *Programs and abstracts, Child Neurological Society*, 26, 483-484.
- Ritvo, E.R., Freeman, B.J., Scheibel, A.B., Duong, P.T., Robinson, H., & Guthrie, D. (1986) Lower Purkinje cell counts in the cerebella of four autistic patients: Initial findings of the UCLA-NSAC autopsy research report. *American Journal of Psychiatry*, 143, 862-866.

- Roberts, S. & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107, 358-367.
- Rumelhart, D.E. & McClelland, J.L. (1986). On learning the past tenses of English verbs. In D.E. Rumelhart, J.L. McClelland & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and Biological Models* (pp. 216-271). Cambridge, MA: MIT Press/Bradford Books.
- Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, J.L. McClelland & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations* (pp. 318-362). Cambridge, MA: MIT Press/Bradford Books.
- Ruppin, E., Reggia, J.A., & Horn, D. (1996). Pathogenesis of schizophrenic delusions and hallucinations: A neural model. *Schizophrenia Bulletin*, 22, 105-123.
- Sartorius, N. Shapiro, R., & Jablonsky, A. (1974). The international pilot study of schizophrenia. *Schizophrenia Bulletin*, 1, 21-35.
- Schreibman, L. & Lovaas, O.I. (1973). Overselective responding to social stimuli by autistic children. *Journal of Abnormal Child Psychology*, 1, 152-168.
- Seidenberg, M.S., Bruck, M., Fornarolo, G., Backman, J. (1985). Word recognition processes of poor and disabled readers: Do they necessarily differ? *Applied Linguistics*, 6, 161-179.
- Seidenberg, M.S., & McClelland, J.L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523-568.

- Selemon, L.D., Rajkowska, G., & Goldman-Rakic, P.S. (1995). Abnormally high neuronal density in the schizophrenic cortex: A morphometric analysis of prefrontal area 9 and occipital area 17. *Archives of General Psychiatry*, 52, 805-818.
- Snowling, M. (1987). *Dyslexia: A cognitive developmental perspective*. Oxford: Basil Blackwell.
- Snowling, M., Hulme, C., & Goulandris, N. (1994). Word recognition in developmental dyslexia: A connectionist interpretation. *The Quarterly Journal of Experimental Psychology*, 47A, 895-916.
- Stanovich, K.E. (1988). Explaining the differences between the dyslexic and the garden-variety poor reader: The phonological-core variable-difference model. *Journal of Learning Disabilities*, 21, 590-604.
- Tallal, P. (1990). Fine-grained discrimination deficits in language-learning impaired children are specific neither to the auditory modality nor to speech perception. *Journal of Speech and Hearing Research*, 33, 616-617.
- Tallal, P., Miller, S.L., Bedi, G., Byrna, G., Wang, X., Srikantan, S.N., Schreiner, C., Jenkins, W.M., & Merzenich, M.M. (1996). Language comprehension in language-learning impaired children improved with acoustically modified speech. *Science*, 271, 81-84.
- Temple, C.M. (1997). Cognitive neuropsychology and its applications to children. *Journal of Child Psychology and Psychiatry*, 38, 27-52.
- Thagard, P. (1992). *Conceptual Revolutions*. Princeton, NJ: Princeton University Press.

- Thomas, M. S. C., Grant, J., Barham, Z., Gsodl, M., Laing, E., Lakusta, L., Tylers, L. K., Grice, S., Paterson, S., & Karmiloff-Smith, A. (2001). Past tense formation in Williams syndrome. *Language and Cognitive Processes, 16* (2/3), 143-176.
- Thomas, M.S.C., & Karmiloff-Smith, A. (2003-a). Modeling language acquisition in atypical phenotypes. *Psychological Review, 110*.
- Thomas, M. & Karmiloff-Smith, A. (2003-b). Connectionist models of development, developmental disorders and individual differences. In R.J. Sternberg, J. Lautrey, & T. Lubart (Eds.), *Models of intelligence: International perspectives* (pp. 133-150). American Psychological Association.
- Thomas, M.S.C. & Karmiloff-Smith, A. (2002). Are developmental disorders like cases of adult brain damage? Implications from connectionist modeling. *Behavioral and Brain Sciences, 25*, 727-788.
- Torgesen, J.K. (1999). Reading disabilities. In R. Gallimore & L.P. Bernheimer (Eds.) *Developmental perspectives on children with high-incidence disabilities: The LEA series on special education and disability* (pp. 157-181). Mahwah, NJ: Lawrence Erlbaum.
- Treiman, R. & Hirsh-Pasek, K. (1985). Are there qualitative differences in reading behavior between dyslexics and normal readers? *Memory and Cognition, 13*, 357-364.
- Van Orden, G.C., Bosman, A.M.T., Goldinger, S.D., & Farrar, W.T. (1997). A recurrent-network account of reading, spelling, and dyslexia. In J.W. Donahoe & V.P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations*. Amsterdam, The Netherlands: Elsevier Science.

- Weinberger, D.R., Berman, K.F., Suddath, R., & Torrey, E.F. (1992). Evidence of dysfunction of a prefrontal-limbic network in schizophrenia: A magnetic resonance imaging and regional cerebral blood flow study of discordant monozygotic twins. *American Journal of Psychiatry*, *149*, 890-897.
- Wellman, H.M., Ritter, K., & Flavell, J.H. (1975). Deliberate memory behavior in the delayed reactions of very young children. *Developmental Psychology*, *11*, 780-787.
- Werker, J., & Tees, R. (1987). Speech perception in severely disabled and average reading children. *Canadian Journal of Psychology*, *41*, 48-61.
- Wimmer, H., & Perner, J. (1983). Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, *13*, 103-128.
- Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1131-1161.

Author Note

Christopher M. Conway, Rick Dale, and Morten H. Christiansen, Department of Psychology; Cornell University; Michelle R. Ellefson, Department of Psychology; University of Warwick.

Correspondence concerning this article should be addressed to: Christopher M. Conway, Department of Psychology, Cornell University, Uris Hall, Room 211, Ithaca, New York, 14853. E-mail: [cmc82@cornell.edu](mailto:cmc82@cornell.edu).

Footnote

<sup>1</sup> Anumolu, Norman, & Bray (1997) do provide a measure of the network's recall accuracy in relation to strategy use; however, such a measure apparently was not applied toward the simulations under discussion.

Table 1

*Connectionist Models of Developmental Disorders*

<u>Group/Model</u>	<u>Disorder</u>	<u>Hypothesized Underlying Factor</u>	<u>Deficit Implementation.</u>
<i>Architecture Manipulations</i>			
Cohen (1998)	autism	high/low neuron density	varied hidden units
Brown & Loosemore (1995)	dyslexia	reduced comp. resources	varied hidden units
Harm & Seidenberg (1999)	dyslexia	phonological impairments	clean-up, units; weight decay
<i>Changes to Training Dynamics</i>			
Bray et al. (1998)	intellectual disability	delayed devel. trajectory	froze weights of normal model
Hoffman & McGlashan (1997)	schizophrenia	atypical brain connectivity	weight pruning
O'Laughlin & Thagard (2000)	autism	weak coherence/integration	modified inhib/excit strength
<i>Changes to I/O Representations</i>			
Hoeffner & McClelland (1993)	SLI	perceptual process deficit	reduced input strength
Brown (1997)	dyslexia	phonological deficit	distributed input representation
Thomas & Karmiloff-Smith (2003-a)	WS	over-emphasis on phonology	increased number I/O units