Proceedings of the 25th Annual Child Language Research Forum. Center for the Study of Language and Information, Stanford University. (1993).

Can a perceptual processing deficit explain the impairment of inflectional morphology in developmental dysphasia? A computational investigation.

James H. Hoeffner and James L. McClelland

Carnegie Mellon University

Abstract: Developmental dysphasia (also referred to as specific language impairment) is a developmental language disorder in which children display delayed or abnormal language development but have normal non-verbal intelligence and no gross perceptual or neurological disorders. Several investigators have hypothesized that the underlying cause of developmental dysphasia is a subtle perceptual processing deficit. But there have also been criticisms of the perceptual processing account (Gopnik & Crago, 1991) and several alternative theories have been put forth (Gopnik & Crago, 1991; Clahsen, 1989). Furthermore, it has recently been argued that the behavior of developmental dysphasics provides strong support for dual mechanism accounts of morphology acquisition and processing (Pinker & Prince, 1991; Gopnik & Crago, 1991).

This paper contributes to the debate by studying the effects of a perceptual processing deficit on a model that simulates the acquisition of morphology. A neural network model that is capable of learning a system of semantics to phonology mappings analogous to part of the system of English inflectional morphology is presented. When the phonological input to the model is selectively distorted in a manner consistent with phonological processing deficits of developmental dysphasics, the model's performance simulates dysphasic children's performance. When compared to the undamaged model, the "dysphasic" model displays slower overall learning, a higher rate of errors on suffixation, differential success on the third person singular, progressive -ing and -ed suffixes, and greater difficulty with regular than irregular past tenses. The model also addresses arguments that have been put forth against the perceptual processing theories, including the assertion that the dysphasics' greater difficulty with morphemic than with non-morphemic phonemes undermines the validity of a perceptual processing account. Gopnik & Crago (1991) argue that a perceptual processing account must predict that dysphasics will have equal difficulty producing identical phonemes regardless of morphological structure. For example, they should have equal difficulty with the -s in nose and the -s in bees. The model demonstrates that this claim is false and that differential performance on identical phonemes can arise from the effects of competition between the members of an inflectional paradigm.

The model demonstrates that many of the symptoms of dysphasic speech can be produced by distorting the phonological input to a unitary learning system. The model's performance provides support for perceptual processing accounts of developmental dysphasia and, like Marchman (1993), calls into question the claim that dysphasics' behavior can be used to support dual mechanism accounts of morphology acquisition.

Introduction: Developmental dysphasia, or specific language impairment (SLI), is a developmental language disorder in which children display delayed or abnormal language development but have normal non-verbal intelligence and no gross perceptual, behavioral, or neurological deficits. Dysphasics form a relatively heterogeneous population and there may in fact be a variety of related disorders all grouped under the rubric of SLI. Nonetheless, many children with SLI have a similar linguistic profile: "...a mild to moderate deficit in a range of language areas and a more serious deficit in the use of morphology." (Leonard et al, 1992, p. 1077)

There have been a variety of theories put forth to explain SLI. One group of theories posits a perceptual processing deficit as the ultimate cause of the children's linguistic difficulties. This hypothesis is based on a large body of work investigating the perceptual processing abilities of children with SLI. A series of investigations carried out by Paula Tallal and her colleagues (Tallal & Piercy, 1973a, 1973b, 1974, 1975; Tallal & Stark, 1981) have shown that dysphasic children suffer from a subtle perceptual processing deficit. This deficit affects their ability to perceive and discriminate rapidly changing sounds. For example, children with SLI are significantly worse than normals at discriminating stop consonant-vowel pairs such as *ba-da*, or fricative-vowel pairs such as *sa-sha*. A perceptual deficit of this kind could plausibly lead to the kinds of linguistic problems evidenced in children with SLI.

Cross-linguistic studies of English, Italian and Hebrew speaking dysphasics conducted by Leonard and his colleagues (Leonard et al, 1988; Leonard, 1989, Leonard, 1993) have added further support for the perceptual processing hypothesis. A comparison of English and Italian dysphasics is particularly instructive. Most of English grammatical morphology is represented by relatively low salience markers (e.g., word final non-syllabic consonants and unstressed syllables). Consequently, English dysphasics have problems with much of their language's grammatical morphology. But most Italian morphological markers are more salient than the English equivalents. In keeping with the perceptual processing hypothesis, Italian children perform much better than English dysphasics (and at levels similar to normal Italian children) on most morphological markers. But the Italian dysphasics are impaired on some morphological markers; those which are low in phonetic substance. A particularly interesting contrast is the fact that the Italian dysphasics perform much better with the vowel-final feminine articles *la* and *una* then they do with the consonant-final masculine equivalents, *il* and *un*. Leonard has hypothesized that the SLI children's difficulties with low salience items may be the consequence of a reduced processing capacity.

Recently, Gopnik & Crago (1991) have proposed that the underlying cause of SLI is neither a perceptual processing deficit nor a more general cognitive problem, but a specifically linguistic deficit. Gopnik & Crago adopt Pinker's (Pinker & Prince, 1991; Marcus et al, 1992), dual mechanism theory of morphology acquisition and processing. According to this theory, the computational demands required for the acquisition and production of regularly and irregularly inflected forms are so different that two separate and qualitatively distinct systems have evolved to handle regular and irregular forms. The regular forms are produced by a rule-based system which functions by appending *-ed* to any verb stem that is not marked as an irregular. The second system is an associative memory system. Through frequent exposure it can learn the irregulars on a case by case basis.

Since these systems are thought to be distinct neurally as well as functionally, we can expect that they may dissociate under certain conditions.

Gopnik & Crago claim that SLI is a result of a genetic deficit which affects the part of the language system responsible for regular inflectional morphology. Because of the defect in the regular rule-based system, children with SLI cannot learn to form morphological paradigms and manipulate grammatical features the way that normals do. Therefore, these children must rely solely on the associative memory system. Since the memory system is performing a task it was not designed for, learning of regular morphology is dramatically impoverished.

In this paper, we contribute to the debate by comparing the behavior of two neural networks as they attempt to learn a system of semantic to phonology mappings. The two networks are identical in all respects but one: the phonological input to the "SLI network" is degraded in a manner consistent with the hypothesized processing deficit of dysphasic children.

Models: The model discussed in this paper is a neural network model that learns to relate semantic and phonological representations. The model was developed to simulate morphology acquisition and processing in normals (Hoeffner, 1992). The architecture of the model is shown in Figure 1. The portion of the figure included in the rectangular box is the implemented model.

Architecture: There are three layers, a semantic layer, a hidden layer and a phonological layer. There are connections between layers, and, for the hidden and phonological layers, there are also bi-directional weights connecting the units within a layer. The bi-directional weights allow the network to perform an iterative computation, passing information back and forth through the network as it gradually settles to a stable state. This type of network is known as an attractor network. There are several reasons for using this type of network: its ability to learn arbitrary mappings, the flexibility regarding input/output relations, the ability to vary input strength (soft-clamping), and the use of settling time as an analog of reaction time (see Hoeffner, 1992, for further explanation).

We envision the current model as being embedded in a larger set of systems (see Figure 1). We assume that a perceptual system is responsible for feeding phonological representations into the morphology learning network. If this system is impaired, then systematically distorted representations will be fed into the phonological layer of the morphology learning system.

Corpus: The training corpus is based on data from the Marcus et al (1992) monograph. It consists of all the monosyllabic regular past tense verbs produced by Adam, Eve, Sarah, Abe, or the adults conversing with them (Table A11 in Marcus et al) and all the irregular past tense verbs produced by the four children (Tables A5-A8).

There are a total of 1925 verb forms in the corpus, arranged in 385 paradigms. Each paradigm has five members: a zero marked form (*jump*), a past tense (*jumped*), third person singular (*jumps*), progressive (*jumping*), and a past participle (*jumped*). The frequencies of each verb are taken from Kucera and Francis (1967).

Phonological representation: The phonology is represented by a feature/slot scheme. There are 8 slots: 3 prevocalic, one vocalic, 3 post vocalic and one syllabic suffix slot. Each consonant slot is made up of ten units representing standard phonological features (e.g., voicing, etc). Consonants spread from the periphery in to cover all three of the pre or post vocalic slots. This representation was chosen in part to avoid a special "suffix" slot, and allow an unconfounded study of the differential effects of the phonological manipulation on morphemic and non-morphemic phonemes.

Figure 1: Architecture of the Models

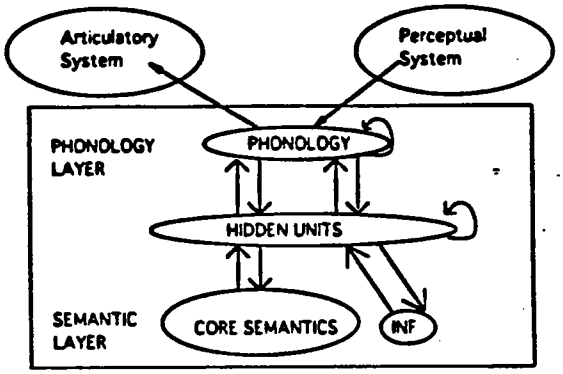

Processing and learning: The networks were trained with the Contrastive Hebbian Learning algorithm (Peterson & Anderson, 1987). Contrastive Hebbian learning makes use of two phases; a plus phase and minus phase. In the plus phase, both semantic and phonological inputs are presented to the network. In the minus phase, only the semantic information is presented. The algorithm adjusts the networks' weights so as to reduce the discrepancy between unit activations in the plus and minus phases. Therefore, a trained network can be given the semantic representation of a word as input and it will be able to output the phonological representation of that word (i.e., create the correct representation across the phonology layer).

Impairing phonology: We simulate SLI by weakening the phonological input to the model. The strength of the phonological input to the normal model was 30 for all phoneme/position combinations except for word-final stops and fricatives, which had a strength of 10. The strengths of all phonemes were reduced by 1/3 in the impaired model. Therefore, word-final stops and fricatives had a strength of 3.3 and all other phoneme/positions had a strength of 10 in the impaired model.

These manipulations are meant to capture three points: 1) some phoneme/position combinations important to morphology acquisition may be of relatively low salience for all learners (hence the word-final stops and fricatives have 1/3 the strength of the other phonemes in both of the models), 2) the strength of the phonological input may be generally weakened for dysphasic learners (so all phonemes are reduced by 1/3), and, 3) although all phonemes may be weaker in the dysphasic case, only the most vulnerable phonemes may drop clearly below the level required for effective learning. The effect of a change from a strength of 30 to a strength of 10 is not great. The model is still able to learn fairly effectively in both cases. But when the strength is dropped as low as 3.3, feedback from the rest of the network can overwhelm the weakened input, resulting in a much noisier and less stable training signal in the impaired case.

Scoring procedure: As shown in Figure 1, the representations across the model's phonology layer have to drive a separate articulatory system. We assume that phonological representations have to be of a sufficient strength or quality to drive articulation. Therefore a strict set of criteria are applied to the model's responses in order to distinguish valid from not valid responses. Only the valid responses are thought to be able to drive articulation. In order to be a valid response, three criteria must be met: each unit must be above +.85 or below -.85.

Figure 2: Normal vs. Impaired Model: Proportion No Responses

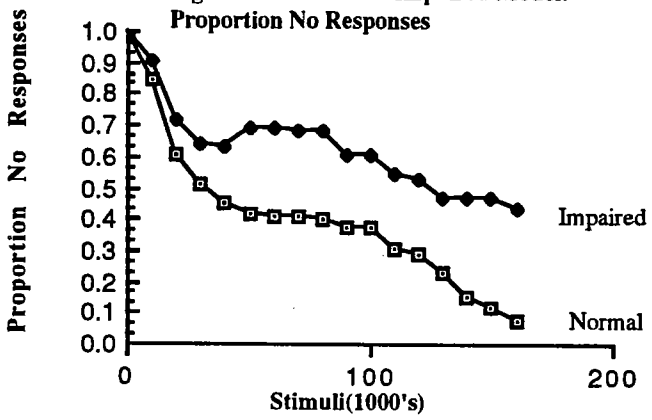

Additionally, each phoneme must be a legal phoneme and the network must settle in less than 100 cycles. If a response does not meet the three criteria then it is classified as a No Response. If a response does meet the three criteria, it is further subcategorized as a Correct Responses or as one of the error types.

Results: The normal model was trained until it reached a preset performance criterion: greater than 90% of the responses were valid responses and greater than 95% of the valid responses were correct. It took the normal model approximately 160,000 stimulus presentations to reach the criteria. The impaired model was trained for the same number of stimulus presentations as the normal model.

Figures 2 and 3 show the models' performance on No Responses and on proportion Correct of Valid Responses. As can be seen in the figures, learning in the impaired model is slower and more error prone than in the normal model. The normal model gradually reduces the number of No Responses as it masters the corpus. By the end of training only 8% of the normal model's responses are classified as No Responses. The normal model also displays a fairly low error rate on its Valid Responses. The model climbs quickly to ~90% correct and maintains performance at close to 90% correct for most of training. By the end of training, the model reached 96% correct.

The impaired model is a much slower learner, even after ~160,000 stimuli presentations it is still producing more than 40% No Responses. And despite the smaller number of valid responses produced by the impaired model, it still makes a larger number of errors on the valid responses than does the normal model, averaging 87% correct overall to the normal model's 93% correct.

But the performance of the impaired model was not just degraded uniformly. Certain of the forms were more adversely affected by the weakening of the phonological input. Figure 4 shows both models' performance on the unmarked (*jump*) and the 3s (*jumps*) forms. There is a slight decrease in performance on unmarked forms, (from 95% to 92.5% correct), but a much greater decrease in performance on the 3s forms (from 91.5% to 63% correct). Performance on the regular past tenses and past participles was also impaired, while performance on the irregular pasts and the -ing forms was relatively spared.

A major reason for the precipitous drop in performance on the 3s forms was a large increase in the number of No Marking Errors (e.g., *jump* is produced instead

Figure 3: Proportion Correct, Normal vs. Impaired Model

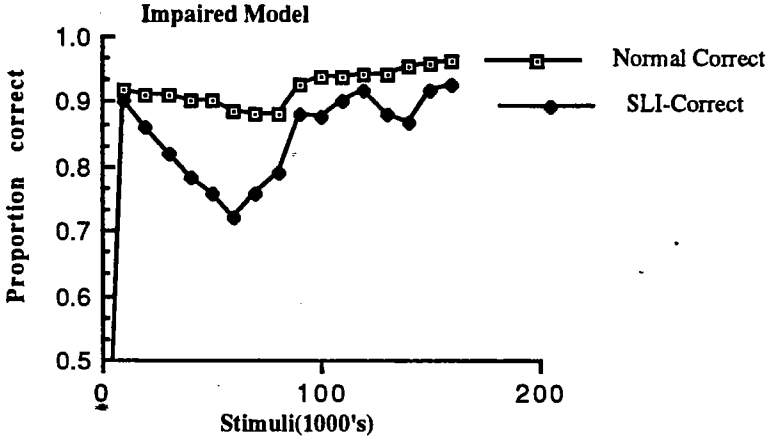

of an inflected form like *jumps* or *jumped*). Both normal children and children with SLI produce No Marking Errors, but children with SLI fail to reliably mark lexical items at a higher rate and for a longer period than do normals. The models replicate these phenomena. The normal model produces some unmarked forms, but the impaired model produces more than double the number of No Marking Errors (see Figure 5).

The impaired model also shows a dissociation between regular and irregular past tenses (see Figure 6). Both regular and irregular pasts are affected by the weakening of the phonological input, but the irregulars are more severely affected. The irregulars dropped from 93% correct to 89% correct, while the regulars dropped from 90% correct in the normal model to only 78% correct in the impaired model. Much of the difference is due to an increase in No Marking Errors for the regular verbs in the impaired model. Although the impaired model's performance (78% correct) on regular pasts is much lower than the normal model's, it is still higher than that seen in many dysphasic children. But this is not a serious problem, since greater decrements in performance could easily be achieved by a further weakening of the phonological input.

Figure 4: Proportion Correct on Unmarked Forms and 3s Forms

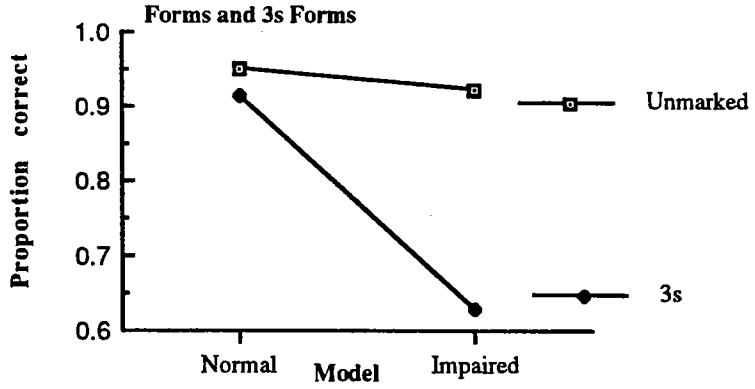

Figure 5: Normal vs Impaired Model: Number of No Marking Errors (All Inflected Forms)

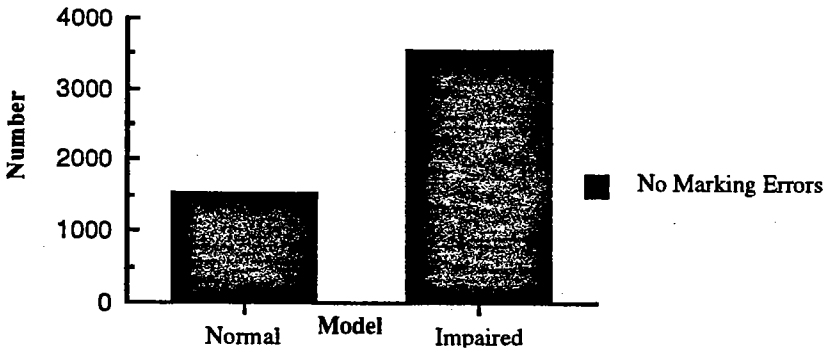

Figure 6: Normal vs Impaired Model: Proportion Correct on Regular and Irregular Pasts

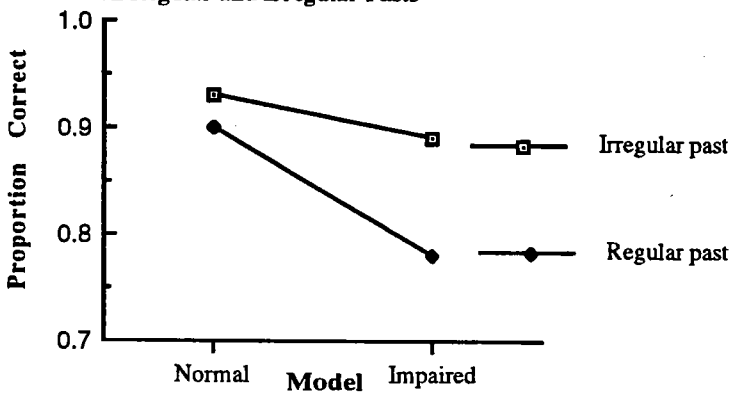

Overregularizations: According to the dual mechanism account of SLI, dysphasic children should not produce overregularized forms of irregular past tenses (*goed, ated,* etc). Gopnik has claimed that dysphasics do not make overregularization errors: "For example, though young children frequently make overregularization errors, such as *foots* or *hitted*, these types of errors do not occur in the dysphasic population. "(Gopnik, 1990, p.141.) But, in fact, dysphasic children do produce overregularization errors both in spontaneous speech and during elicitation tasks (Clahsen, 1989; Leonard, 1993).

Both the normal model and the impaired model produced overregularized forms of irregular pasts. Of course, we expect the normal model to overregularize. But both the impaired model and dysphasic children fail to mark tense on many regular verbs. Why would they simultaneously produce overregularized forms?

The explanation lies in an understanding of the single mechanism account of overregularization. Overregularizations arise as a consequence of a single mechanism being called upon to simultaneously perform several competing tasks. For regular verbs, the network must learn to use the activation of the past tense semantic units to activate the proper suffix while leaving the rest of the word unchanged. For no change irregulars, (*hit, put, cut,* etc.) it must prevent

activation of the suffix and leave the stem unaltered. For other irregulars, it must also prevent activation of the suffix but at the same time it has to perform various semi-regular (*sing-sang*, etc.) or idiosyncratic changes (*go-went*) to the stem. All these different and contradictory mappings must be performed by the same weight matrix. If the network makes an error on one type of mapping, then the learning algorithm will make changes to the weights in order to reduce the error for that verb. But these weight changes may then increase the probability of performing the wrong mapping for another verb. The inevitable result of a single mechanism trying to master such a contradictory system is that there will be times when one mapping interferes with the others. This will be true for dysphasic learners as well as for normals, since, in both cases, a single learning mechanism is being called upon to master a complex, contradictory set of mappings.

Phonologically identical morphemic and non-morphemic items: An important criticism of the perceptual processing account of SLI is that perceptually identical items such as the /z/ in bees and the /z/ in buzz are not similarly affected in SLI. The children are much more likely to drop the morphemic /z/ in *bees* than the non-morphemic equivalent in *buzz*. They also display different rates of errors on phonologically identical morphemes such as the plural *-s* and the the third person *-s*. Gopnik & Crago have argued that for the perceptual processing account to be correct: "all similar surface forms of the utterance must be affected in a similar way." (Gopnik & Crago, 1991).

But there are several problems with Gopnik & Crago's claim. The perceptual processing theories do not need to claim that perceptual factors are the only relevant variable. It is likely that perceptual factors interact with a variety of other factors such as frequency, morphological structure, and semantic factors. We know that normal children acquiring English do not acquire homophonous morphemes such as the plural -s, third person -s, and possessive -s at identical rates (Brown, 1973). There is no reason for us to assume that the same factors that determine the differential acquisition of these morphemes in normal children may not also be operating in the dysphasic case. If perceptual salience interacts with these factors then we would expect that phonologically identical morphemes could be differentially affected.

The model's performance supports this claim. If we compare the impaired model's performance on words like *add* and *need* that have non-morphemic final *-d* 's to its performance on words like *tried* and *died*, we see that the forms are differentially affected despite the fact that they underwent identical phonological weakening. *Add* was produced correctly 96% of the time and *need* 100%, while *tried* was correct only 77% and *died* only 10% of the time. But the differences are not just in the relative error rates, but more importantly, in the types of errors committed. For the two morphemic final -d forms (*tried*, *died*) there were 41 No Marking Errors. But, for the non-morphemic final -d forms, *need* and *add*, there were none. In fact, on all the forms in the corpus that end in a non-morphemic final -d there is not one example of the equivalent of a No Marking Error (i.e., producing *nee* for *need*). The reason is simple: there is no form in the corpus that is closely related to *need* semantically and also has the form *nee*. Therefore, *need* doesn't have to compete with *nee*, but *died* does have to compete with *die*. Of course, this intra-paradigm competition takes place in the normal model as well, but in the SLI model the phonological impairment shifts the balance in the competition in favor of the unmarked form.

The main effects of the phonological manipulation on the model's performance can be summarized in six points:

1) The impaired model displays generally slower and more error prone learning.

2) Past tense, past participle and 3s forms show greater impairment than the unmarked forms or the -ing forms.

3) There is a dramatic increase in No Marking Errors on all the affected inflected forms, just as we see in the case of English dysphasics.
4) Regular pasts are affected more than irregular pasts.
5) Both models occasionally overregularize irregular pasts.
6) Morphemic and non-morphemic phonemes can be differentially affected even if they are phonologically identical. One of the reasons may be an interaction between perceptual factors and intra-paradigm competition.

Conclusions: Since many of the most salient features of dysphasic speech can be produced by weakening the phonological input to a "normal" learning mechanism, the model's performance provides further support for theories that posit a perceptual deficit as the underlying cause of SLI. The fact that phonologically identical morphemic and non-morphemic phonemes can be treated differently by the impaired model argues against one of the primary objections to the perceptual deficit theories. Furthermore, the dissociation between regular and irregular past tense forms that was produced by impairing the phonological input to a unitary learning mechanism argues against the dual mechanism interpretation of SLI (Marchman, 1993). Additionally, the fact that children with SLI occasionally produce overregularizations of irregular pasts also argues against the dual mechanism theory, but is compatible with the single mechanism account.

Although the model presented here was able to simulate many aspects of SLI children's acquisition of English morphology, there are some limitations to the current work. We know, for example, that the linguistic problems of children with SLI are not restricted to morphophonology; they also often have syntactic difficulties. It may be that at least some of the SLI children's syntactic difficulties can arise more or less directly from their problems with morphology. For example, the difference between passive and active constructions in English is signalled by a few relatively low salience morphemes (*was*, *-ed*, *by*). Difficulties with these lower salience items could interfere with the learning of active and passive constructions. It is also possible to extend the present model to simulate the acquisition of more than main verb bound morphology. For example, it would be fairly simple to also represent nominal morphology or to add units to the phonological and semantic layers to represent modals and auxiliaries as well as main verbs.

Another important issue to address in future work is the effects of capacity limitations on the models' performance. In the current model we assumed that a perceptual processing deficit resulted in distorted representations being fed into an otherwise normal learning system. We did not specify what the ultimate cause of the perceptual deficit is, but the model is consistent with Leonard's hypothesis of a reduced processing capacity that leads to a differential degradation of the more difficult (i.e., less salient) items. But there is also another notion of capacity that may also be relevant to an understanding of SLI. This notion of capacity has less to do with on-line processing and more to do with the overall resources available to the learner. In connectionist models, limitations in learning resources can be simulated by restricting the number of hidden units or the overall number of weights available to the network. Recently, Marchman (1993) showed that a dissociation between regular and irregular verbs could be produced by restricting the learning resources available to a feed-forward verb learning network. In the future, it will be important to explore the effects of capacity limitations on the performance of the networks discussed in this paper. It may be that both a perceptual impairment based on reduced on-line capacity and limitations in overall learning resources can independently produce behavior symptomatic of

SLI, or it may be that the combination of both a perceptual impairment and resource restrictions will provide the best explanation of dysphasic behavior.

References

Brown, R. (1973) *A first language*. Cambridge, MA: Harvard University Press.

Clahsen, H. (1989). The grammatical characterization of developmental dysphasia. *Linguistics*, 27, 897-920.

Gopnik, M. (1990) Feature-blindness: A case study. *Language Acquisition*, 1, 139-164.

Gopnik M., & Crago, M. (1991) Familial aggregation of a developmental language disorder. *Cognition*, 39, 1-50.

Hoeffner, J. (1992). Are rules a thing of the past? The acquisition of verbal morphology by an attractor network. *In Proceedings of the 14th Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Kucera, N. & Francis, W.N., (1967). *Computational analysis of present day American English*. Providence RI: Brown University Press.

Leonard, L. (1989) Language learnability and specific language impairment in children. *Applied Psycholinguistics*, 10, 179-202.

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., 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.

Leonard, L., McGregor, K., Allen, G. (1992). Grammatical morphology and speech perception in children with specific language impairment. *Journal of Speech and Hearing Research*, 35, 1076-1085.

Marchman, V., A. (1992). Constraints on plasticity in a connectionist model of the English past tense. *Journal of Cognitive Neuroscience*, 5:2 215-34.

Marcus, G.F., Pinker, S., Ullman, M., Hollander, M., Rosen, T.J., Xu, F. (1992) Overregularization in language acquisition. SRCD Monograph, Serial No. 228, Vol. 57, No. 4.

Peterson C. & Anderson, J.R. (1987) A Mean Field Theory Learning Algorithm for Neural Networks. *Complex Systems* 1, 995-1019.

Pinker, S., & Prince, A. (1991) Regular and irregular morphology and the psychological status of rules of grammar. *Proceedings of the 1991 Meeting of the Berkeley Linguistics Society*.

Tallal, P. and Piercy, M. (1973a). Defects of non-verbal auditory perception in children with developmental aphasia. *Nature* 241, 468-469.

Tallal, P. and Piercy, M. (1973b). Developmental aphasia: Impaired rate of non-verbal processing as a function of sensory modality. *Neuropsychologia*, 11, 389-398.

Tallal, P. and Piercy, M. (1974). Developmental aphasia: rate of auditory processing and selective impairment of consonant perception. *Neuropsychologia* 12, 83-93.

Tallal, P. & Piercy, M. (1975). Developmental Aphasia: The perception of brief vowels and extended stop consonants. *Neuropsychologia*, 13, 69-74.

Tallal, P. & Stark, R. (1981). Speech acoustic cue discrimination abilities of normally developing and language impaired children. *Journal of the Acoustical Society of America*, 69, 568-574.