MODELLING VOCABULARY ACQUISITION: AN EXPLANATION OF THE LINK BETWEEN THE PHONOLOGICAL LOOP AND LONG-TERM MEMORY

GARY JONES
FERNAND GOBET
JULIAN M. PINE
ESRC Centre for Research in Development, Instruction and Training
Department of Psychology; University of Nottingham,
Nottingham, NG7 2RD, England

ESRC Centre for Research in Development, Instruction and Training
May 1999
Technical Report No. 61

Abstract

The acquisition of vocabulary represents a key phenomenon in language acquisition, but is still poorly understood. Recently, the working memory model (Baddeley & Hitch, 1974) has been adapted to account for vocabulary acquisition (e.g. Gathercole & Baddeley, 1989). It is claimed that the phonological store, one of the components of working memory, offers a critical mechanism for learning new words. One of the theoretical weaknesses of this approach is however that no account is given for the mechanisms and representations used in long-term memory learning. This paper presents a computer model combining the EPAM/chunking approach (Feigenbaum & Simon, 1984) with the working memory approach. Phonemic learning is simulated as the elaboration of a discrimination net. Naturalistic input, consisting of utterances from five mothers interacting with their child, is used during the learning phase. Simulations show that the model can account reasonably well for the nonword repetition task described by Gathercole and Baddeley (1989), a task often presented as a powerful diagnostic of vocabulary learning.

E-mail: {gaj, frg, jp}@Psychology.Nottingham.AC.UK

1 INTRODUCTION

Children are remarkably adept at learning new verbal information. After an initial slow period from about 12 to 16 months when most children learn around 40 words, the learning rate increases such that in the next four months children will have learnt 130 more new words (Bates et al., 1994). By the beginning of their school years children learn up to 3000 words per year (Nagy & Herman, 1987).

A major part of learning new words is learning the novel sequences of sounds that represent the word. However, it is difficult to directly examine the processes involved in learning the sound patterns of new words because it is impossible to be certain that the new sound pattern has never been encountered before. The use of nonwords which conform to the phonotactic rules of English provides a good test of vocabulary learning because it ensures that the (non)word to be learned is novel.

1.1 The Nonword Repetition Test

The nonword repetition (NWR) test (Gathercole, Willis, Baddeley & Emslie, 1994) was designed to investigate the role of phonological memory in word learning. The test involves the experimenter speaking a nonword and the child attempting to repeat the nonword. Gathercole and Baddeley (1989) found that the NWR test was the best predictor of vocabulary size, even after vocabulary scores (as calculated by the British Picture Vocabulary Scale, Dunn & Dunn, 1982) are partialled out of correlations.

Gathercole and Baddeley (1990) used the NWR test to categorise children into two groups, those with low NWR scores, and those with high NWR scores. Children in the high NWR group were better at learning nonword labels than children in the low NWR group. Gathercole, Willis, Emslie, and Baddeley (1991) found better NWR performance on nonwords that were rated high in wordlikeness than nonwords rated low in wordlikeness. These nonword repetition studies show the influence that vocabulary knowledge has upon the learning of new words.

The nonword repetition test involves two sets of nonwords, one set having single consonants (e.g. rubid) and one set having clustered consonants (e.g. glistow). There are twenty nonwords in each set, divided into four groups of five based on the number of syllables in the nonword (one to four). Several studies using these types of nonwords have consistently found that repetition accuracy decreases as the number of syllables in the nonword increases, excepting one-syllable nonwords (e.g. Gathercole & Adams, 1993; Gathercole, Willis, Emslie & Baddeley, 1991), and that accuracy is worse for clustered consonant nonwords.

Based on these findings, nonword repetition ability would seem to provide a good test of phonological memory and is a good indicator of vocabulary size. The NWR findings can be explained within the theoretical framework of the working memory model.

1.2 The Phonological Loop Explanation of Nonword Repetition Findings

The working memory model (Baddeley & Hitch, 1974) has recently been adapted to account for vocabulary acquisition (e.g. Gathercole & Baddeley, 1989). It is claimed that the phonological loop part of the model is a critical mechanism for learning new words. The phonological loop has two linked components: the phonological short-term store, and the sub-vocal rehearsal mechanism. Items in the store decay over time (around 2000 ms, Baddeley, Thomson & Buchanan, 1975). The sub-vocal rehearsal mechanism (involving sub-vocal articulation in real-time) can refresh items in the store in a serial, time-based manner (Gathercole & Martin, 1996). The store is linked to the central executive part of the model, which provides a link to long-term memory (Gathercole & Baddeley, 1993). The influence of long-term memory is acknowledged (e.g. Gathercole, Willis, Baddeley & Emslie, 1994), but the nature and definition of the link is yet to be defined.

The phonological loop hypothesis is able to explain the basic nonword repetition findings involving nonword length because of the decay that takes place in the phonological store (items remain in the store for ≈2000 ms unless refreshed). Longer nonwords take longer to rehearse and so their representations in the phonological store are not refreshed as often as shorter nonwords. Repetition ability will therefore decrease for longer nonwords. The poorer repetition performance for clustered consonant nonwords can be explained in a similar way (because although the clustered consonant nonwords contain the same number of syllables as the single consonant nonwords, they contain more phonemes).

Differences in nonword repetition ability between children of the same age were originally attributed to differences in rates of subvocal rehearsal (Gathercole & Baddeley, 1993). However, recent findings show that children do not use subvocal rehearsal until around seven years of age (see Cowan & Kail, 1996, for a review). This has led to the phonological store being assumed to be the primary language learning device, with differences in language learning across children of the same age being explained by the quality of the phonological representation of just-spoken items (Baddeley, Gathercole & Papagno, 1998, p.168). The lack of a rehearsal process for children below seven years of age does not appear to hinder the model. Brown and Hulme (1996) show that nonword repetition phenomena can be explained solely by a decay based model (which the phonological store now becomes for children under seven years of age).

The phonological loop is able to explain a lot of the vocabulary acquisition phenomena using a very simple mechanism. However, it fails in two critical areas: there is no explanation of how words are learned, and there is no explanation of how the loop interacts with long-term memory. Speculative explanations have been given as to how the loop may interact with long-term memory (e.g. Baddeley, Gathercole & Papagno, 1998; Gathercole & Baddeley, 1989). Gupta and MacWhinney (1997) have proposed a formal specification but their model has not yet been implemented computationally.

The phonological loop hypothesis also lacks precision because it is part of a verbal theory. For example, there is no definition of how rehearsal rate changes based on the length of the sound pattern being rehearsed (except to say that long strings are rehearsed slower than short strings). Implementing the loop within a computational architecture forces precision because the theory is implemented as a running computer program. Several computational implementations of the phonological loop exist.

1.3 Existing Computational Implementations of the Phonological Loop

Burgess and Hitch (1992) detail a connectionist network which was primarily intended to model serial order effects (the recall of a set of items in the correct presentation order). The decay in the phonological store is represented by decay on the weights between layer nodes (decay is proportional to the number of phonemes to be output). Rehearsal is synonymous with articulation: the most active word in the network is selected for output once the phonemic input has been processed. The model can explain word length and articulatory suppression effects, but does not explain any of the NWR findings. In addition, no phoneme or word learning takes place; the model provides no theory as to how phonemes and words are created in long-term memory.

Brown and Hulme (1996) give an account of a trace-based decay model which bears resemblance to the phonological store (the model intentionally has no rehearsal process). Time is represented in 0.1-sec slices; input items are split into segments such that each segment corresponds to a time slice. Longer words therefore take up more segments and so occupy more time slices. As each segment of an input item enters the store, it is given a fixed initial strength, which decays over time. This means that the early *segments* of an input item suffer more decay than the later segments, as well as earlier input *items* suffering more decay than later input items.

The probability of recalling an item is the product of the current strength of each of the segments of the item. The probability is increased to reflect the influence of long-term memory during recall; wordlike nonwords would therefore have their probability increased more than non-wordlike nonwords. Using these mechanisms, the model can account for recall effects for different lengths of both words and nonwords (Brown & Hulme, 1996, p.141). However, the model does not account for any learning processes.

There are problems with each of the existing phonological loop models. For the phonological loop to successfully provide an account of vocabulary learning, a precise specification of its interaction with long-term memory is required. EPAM is a computational modelling architecture which is able to provide such a specification.

2 IMPLEMENTING THE PHONOLOGICAL LOOP WITHIN THE EPAM ARCHITECTURE

EPAM has been successful in simulating various areas of human cognition, such as learning, memory, and perception in chess (De Groot & Gobet, 1996; Gobet, 1993; Simon & Gilmartin, 1973), verbal learning behaviour (Feigenbaum & Simon, 1984), the digit-span task (Richman, Staszewski & Simon, 1995), and the context effect in letter perception (Richman & Simon, 1989). EPAM therefore provides a modelling environment which is well suited for describing how sound patterns can be learnt, and how the phonological loop may be connected to sound patterns in long-term memory.

2.1 The EPAM Architecture

EPAM learns by building a discrimination network. The discrimination network is a hierarchical representation of the input which consists of nodes which are connected to one another by links. Nodes contain information and links between nodes contain tests which must be fulfilled before they can be traversed. For the purposes of modelling the learning of sound patterns, EPAM has been simplified and is henceforth referred to as EPAM-VOC.

When an input is given to the network, EPAM-VOC traverses down the hierarchy as far as possible. This is done by starting at the top node (the root node) and selecting the first link whose test is fulfilled by the first part of the input. The node at the end of the link now becomes the top node and the rest of the input is applied to all the links below this node. When a node is reached where no further traversing can be done (because the input fulfils none of the tests of the nodes links, or the node is a leaf node), EPAM-VOC compares the information at the node with the input information. Learning now occurs in two ways.

1. Discrimination. When the input information mismatches the information given at the node (the *image*), a new link (i.e. test) and node are added to the tree below the node that has just been reached. The new test will relate to the mismatched part of the input.
2. Familiarisation. When the input information is under-represented by the information at the node, new features (from the input) are added to the information in the node (the image). In EPAM-VOC, the image of a node will always consist of all the information in the links that lead to the node.

Discrimination therefore creates nodes and links, and familiarisation creates or modifies the information contained in nodes. Examples of the discrimination and familiarisation learning mechanisms will be given later.

2.2 Learning Sound Patterns in EPAM-VOC

EPAM-VOC provides an efficient method for representing items in long-term memory. The basic idea is to give as input to the model utterances from mothers speech so that it can learn phonemes and combinations of phonemes. The mothers' utterances will be converted into a sequence of phonemes before being used as input. This will be done using the CMU Lexicon database (available at http://www.speech.cs.cmu.edu/cgi-bin/cmudict) which cross-references words with their phonemic representations. The use of phonemic input assumes that some form of phonemic feature primitives already exist to distinguish one phoneme from another.

EPAM-VOC will begin with a null root node. When it sees an input (a sequence of phonemes), new nodes and links will be created. At first, most of the new nodes and links will just be for single phonemes, as EPAM-VOC learns to master individual phonemes. As learning progresses, the information at nodes will become sequences of phonemes and therefore segments of speech (e.g. specific words) rather than just individual sounds (i.e. phonemes). To accomplish this, the EPAM learning mechanism is altered in two ways. First, before a sequence of phonemes can be learnt, the individual phonemes in the sequence must have been learnt. Second, when individual phonemes are learnt, they are linked to the root node (in this way all sequences of phonemes are below the node which represents the initial phoneme in the sequence).

Let us consider an example of the network learning the utterance "What about that". Using the CMU Lexicon database, this utterance is converted to the phonemic representation "W AH1 T AH0 B AW1 T DH AE1 T" (all of the phonemes used in the database map onto the standard phoneme set for American English). Note that the phonemic input to the model does not specify gaps between words, but does specify the stress of particular phonemes (0=unstressed; 1=primary stress; 2=secondary stress).

The first part of the input ("W") is applied to all of the root nodes' links in the network. If the network is empty, there will be no links. At this point EPAM-VOC must discriminate because the information "W" mismatches the information at the root node (the root node information is null). The discrimination process creates a new node, and a link from the root node to the new node with the test "W". EPAM-VOC must then familiarise itself with the input, in order to create the "W" information in the image of the node. When the input is seen for the first time, just the phoneme "W" is learnt.

When encountering the input for the second time, the link "W" can be taken, and the input can move to the next phoneme, "AH1". As node "W" does not have any links, then under normal circumstances discrimination would occur below the "W" node. However, the model has not yet learnt the phoneme "AH1", and so discrimination occurs below the root node. Familiarisation will then fill the image of the new node with "AH1".

The third time the input is seen, the "W" link can be taken, with the input moving on to "AH1". No further links are available, but the input "W AH1" does not match the information at node "W" and so discrimination occurs. A new node "W AH1" is created. Familiarisation will fill in the image of the new node. After two further presentations of the input, the network is as shown in Figure 1. The simple example serves to illustrate how EPAM-VOC works; in the actual learning phase each utterance line is only used once, encouraging a diverse network of nodes to be built.

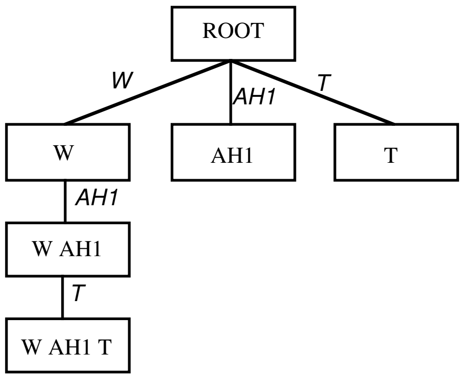

Figure 1: Structure of an EPAM-VOC net after seeing the input "W AH1 T AH0 B AW1 T DH AE1 T" five times.

2.3 Implementing the Phonological Loop and Linking it to Long-Term Memory

The model now requires a specification of the phonological loop and a method by which the loop interacts with long-term memory. The findings relating to the nonword repetition test (the standard test of the phonological loop) were all carried out on children below the age of six. As children below the age of seven are believed not to rehearse, the rehearsal part of the loop should not be used to simulate the nonword repetition findings reported above.

The storage part of the phonological loop is a decay based store which allows items to remain in the store for ≈2000 ms. The model will have a time-limited store which will allow 2000 ms of input (i.e. consistent with the phonological loop estimates). The input will be cut-off as soon as the time limit is reached, because there is no rehearsal to refresh the input representations.

The cumulative time required by the input provides a theory of how the amount of information in the phonological store is mediated by long-term memory. When an input is heard, long-term memory (the EPAM-VOC network) is accessed and the input is represented using the minimum number of nodes possible. Rather than the actual input being placed in the phonological store, the nodes which capture the input are used. The length of time taken to represent the input is therefore calculated on the number of nodes that are required to represent the input. The time allocations are based on estimates from Zhang and Simon (1985), who estimate ≈400 ms to match each node, and ≈84 ms to match each syllable in a node except the first (which takes 0 ms). As the input will be in terms of phonemes, with approximately 2.8 phonemes per syllable (based on estimates from the nonwords in the nonword repetition test), the time to match each phoneme in a node is 30 ms.

Using the example input "W AH1 T AH0 B AW1 T DH AE1 T" and the network as given in Figure 1, the actual input to the model will be "W AH1 T AH0 B AW1". The "W AH1 T" part of the input is represented by a single node, and is allocated a time of 460 ms. Most of the other phonemes are not known to the model and are therefore assumed to take the same time as a full node (400 ms).[1] The actual input to the model therefore has a time allocation of 1660 ms. When the EPAM-VOC network is small, only a small amount of the input information can be represented in the store, and so new nodes will not contain much information. When the EPAM-VOC network is large, a lot of the input information can be represented in the store and so the model can create new nodes which contain large amounts of information.

3 SIMULATING THE NONWORD REPETITION RESULTS

There are two main sets of results for the nonword repetition test. One set was tested on children of four and five years of age (Gathercole & Baddeley, 1989; these results were reported in the introduction). The problem with this study is that the children are of an age where they already have a reasonably large vocabulary size. A second set of nonword repetition results is reported by Gathercole and Adams (1993), who used a simpler version of the test on children of two and three years of age. They found the modified test still allowed phonological memory skills to be reliably tested.

The EPAM-VOC model is able to learn sequences of sounds from inputs that are strings of phonemes (converted from speech utterances). It should therefore be capable of simulating both sets of nonword repetition data (2-3 year olds and 4-5 year olds) by modifying the input that is given to the model to reflect the type of input that will be received by these age groups. *No mechanism within the model will be changed - the only difference between the simulations will be the input given to the model.*

In the simulations that will be presented, the model normally learns something for every input it receives. The EPAM parameter of 8 s to learn a single node was dispensed with so that simulation runs were quicker in real-time. Given the same input to the model, there should be no significant difference to the results whatever the time taken to learn a node.

The NWR test for the model will be performed by presenting each nonword to the model (as a string of phonemes), and seeing if the components of the nonword can be accessed within the same time limitations that were used for the input (see above). By definition, the information at one node will not be able to represent all of a nonword (because the nonword will never have been seen as input). The information from several different nodes will be required to represent the nonword. If the number of nodes and the phonemes in each node can fit into the time limit, the nonword is repeated accurately, otherwise the nonword is repeated incorrectly. The models' NWR test does not involve articulation because the current EPAM-VOC model does not include a theory of articulation.

[1] The time allocated to each phoneme is assumed to be constant.

3.1 Simulation of Two to Three Year Old Children

For the simulation of the NWR test for children of 2-3 years of age, naturalistic input was used for the model. The input consisted of the mother utterances from five mothers interacting with their 2-3 year old children, taken from Lieven and Pine (1995). The average number of utterances for each mother was 27172 (range 23583-32558).

The model was run once for each of the mothers input. After every 500 utterances seen by the model, a nonword test was carried out. This consisted of presenting each nonword as input to the model and seeing if it could represent the nonword within the 2000 ms time capacity. Note that Gathercole and Adams used a simplified version of the NWR test (using 1-3 syllable nonwords and not distinguishing between single and clustered consonants). They also performed a word repetition test. The results of the children and the model (after 4000 utterances of input) for both the nonword and word repetition tests are shown in Figures 2 and 3.

The model performs at ceiling for the one and two-syllable nonwords and words. This is because a minimum of five phonemes can fit into the 2000 ms time limit in the phonological store (one phoneme uses up 400 ms at most). In the nonwords and words used, the average number of phonemes for one-syllable items is 3.2 and for two-syllable items is 5.0. The children do not perform at ceiling for any of the conditions.

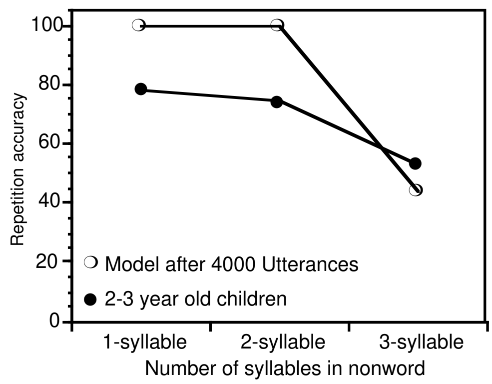

Figure 2: NWR accuracy for 2-3 year old children and the model.

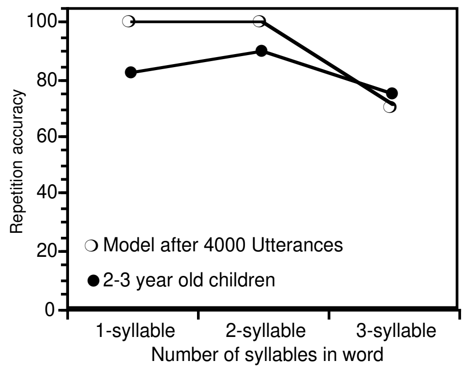

Figure 3: Word repetition accuracy for 2-3 year old children and the model.

The children may not perform at ceiling for one and two-syllable items because of noise during either recognition or articulation of the item. In terms of the model, the children may not perform at ceiling because the phonological store of young children has less duration than adults (all of the timing estimates used in the model are based on studies involving adults).

The results show that the EPAM-VOC model can produce repetition results which are comparable to young children using a simple learning mechanism and naturalistic input. In particular, repetition performance for three syllable items is closely matched by the model. The simulation also raises further questions about children's performance on repetition tests for one and two-syllable items.

3.2 Simulation of Four to Five Year Old Children

Two different types of input were used to simulate the NWR results of 4 and 5 year old children. The same utterances used as input to the model in the simulation of 2-3 year old children formed one set of input. Using only utterances from mothers is not expected to provide a representative sample of the input that 4 and 5 year olds receive, however, because by this age the children are beginning school, and beginning to read. A second input set consisting of a random selection of words from the CMU Lexicon database was therefore used. Figures 4 and 5 show the comparisons of the results of the simulations and the children, for single and clustered nonwords.

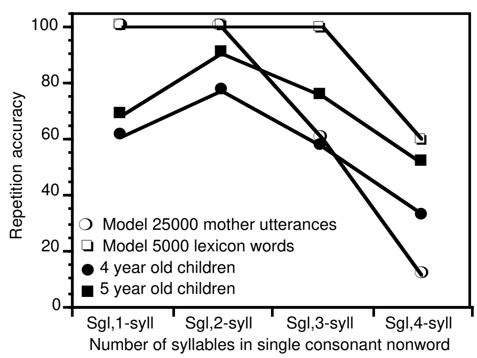

Figure 4: Single consonant NWR accuracy for 4 and 5 year old children, and the model.

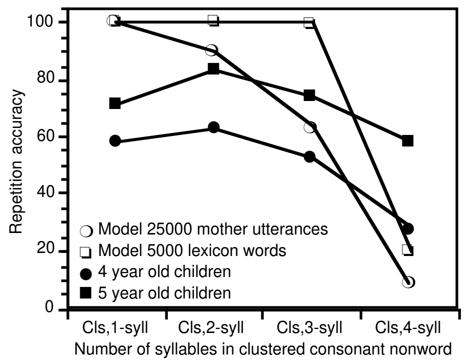

Figure 5: Clustered consonant NWR accuracy for 4 and 5 year old children, and the model.

The children's performance for one-syllable nonwords is actually worse than for two-syllable nonwords (which should be more difficult). The poor performance for children's repetition of one-syllable nonwords is partly due to the acoustic characteristics of their monosyllabic stimuli (Gathercole & Baddeley, 1989).

The model again performs at ceiling for one and two-syllable nonwords (bar one exception). However, the three and four-syllable performance by the model is very interesting. The model tends to over-perform when the lexicon is used as input, yet tends to under-perform when the mothers utterances are used as input. *This highlights the role of the input as a mediating factor in repetition performance.* Even after the model receives 25000 of the mothers utterances as input, its performance is still poor (this is because the input is not diverse in its vocabulary; only 3046 words on average are used in the mother utterances in each of the mother-child interactions). For more diverse input (the words from the lexicon) only 5000 words, each word being treated as a separate utterance, are required for the model to over-perform. The difference in performance across different inputs shows that the input which a child receives will influence the child's repetition performance. The results suggest that a more realistic input would produce a good match to the data.

4 DISCUSSION

The simulations have shown that the EPAM-VOC model is able to approximate the NWR performance of both 2-3 and 4-5 year old children. The model accomplishes this by using a combination of a simple learning mechanism, naturalistic input, and a simple implementation of the phonological loop. This represents a parsimonious approach to learning novel sound patterns. In particular, the model is able to give a specific account of how existing vocabulary knowledge influences the learning of new sound sequences.

The EPAM learning mechanisms are very sensitive to the input that the model receives. This allows the model to make very precise predictions. For example, the sound patterns of the most frequent words in the input will be learnt first. New words which consist of frequent sound patterns will be learned quicker than new words which consist of infrequent sound patterns.

The model also allows comparisons of how different approaches to memory can affect learning. A time based store as presented here can be compared to a chunk based store (e.g. Miller, 1956) and the effects on learning can be examined.

This work represents the first steps in a modelling research program which aims to examine the extent to which the linguistic input a child receives can account for the childs' language development.

REFERENCES

Baddeley, A., Gathercole, S., & Papagno, C. (1998). The phonological loop as a language learning device. *Psychological Review*, *105,* 158-173.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. Bower (Ed.)*, The psychology of learning and motivation*, (Vol. 8), 47-90. New York: Academic Press.

Baddeley, A. D., Thomson, N., & Buchanan, M. (1975). Word length and the structure of short-term memory. *Journal of Verbal Learning and Verbal Behavior*, *14,* 575-589.

Bates, E., Marchman, V., Thal, D., Fenson, L., Dale, P., Reznick, J. S., Reilly, J., & Hartung, J. (1994). Developmental and stylistic variation in the composition of early vocabulary. *Journal of Child Language*, *21,* 85-123.

Brown, G. D. A., & Hulme, C. (1996). Nonword repetition, STM, and word age-of-acquisition: A computational model. In S. E. Gathercole (Ed.)*, Models of short-term memory*, 129-148. Hove, UK: Psychology Press.

Burgess, N., & Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language*, *31,* 429-460.

Cowan, N., & Kail, R. (1996). Covert processes and their development in short-term memory. In S. E. Gathercole (Ed.)*, Models of short-term memory*, 29-50. Hove, UK: Psychology Press.

De Groot, A. D., & Gobet, F. (1996). *Perception and memory in chess: Studies in the heuristics of the professional eye*. Assen: Van Gorcum.

Dunn, L. M., & Dunn, L. M. (1982). *British Picture Vocabulary Scale*. Windsor: NFER-Nelson.

Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, *8,* 305-336.

Gathercole, S. E. (1997). Models of short-term memory. In M. A. Conway (Ed.)*, Cognitive models of memory*, 13-45. Hove, UK: Psychology Press.

Gathercole, S. E., & Adams, A.-M. (1993). Phonological working memory in very young children. *Developmental Psychology*, *29,* 770-778.

Gathercole, S. E., & Baddeley, A. D. (1989). Evaluation of the role of phonological STM in the development of vocabulary in children: A longitudinal study. *Journal of Memory and Language*, *28,* 200-213.

Gathercole, S. E., & Baddeley, A. D. (1990). The role of phonological memory in vocabulary acquisition: A study of young children learning new names. *British Journal of Psychology*, *81,* 439-454.

Gathercole, S. E., & Baddeley, A. D. (1993). *Working memory and language*. Hove, UK: Lawrence Erlbaum Associates.

Gathercole, S. E., & Martin, A. J. (1996). Interactive processes in phonological memory. In S. E. Gathercole (Ed.)*, Models of short-term memory*, 73-100. Hove, UK: Psychology Press.

Gathercole, S. E., Willis, C., Emslie, H., & Baddeley, A. D. (1991). The influences of number of syllables and wordlikeness on children's repetition of nonwords. *Applied Psycholinguistics*, *12,* 349-367.

Gathercole, S. E., Willis, C. S., Baddeley, A. D., & Emslie, H. (1994). The children's test of nonword repetition: A test of phonological working memory. *Memory*, *2,* 103-127.

Gobet, F. (1993). A computer model of chess memory. In *Proceedings of the Fifteenth Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gupta, P., & MacWhinney, B. (1997). Vocabulary acquisition and verbal short-term memory: Computational and neural bases. *Brain and Language*, *59,* 267-333.

Miller, G. A. (1956). The magic number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63,* 81-93.

Nagy, W. E., & Herman, P. A. (1987). Breadth and depth of vocabulary knowledge: Implications for acquisition and instruction. In M. G. McKeown & M. E. Curtis (Eds.), *The nature of vocabulary acquisition*, 19-35. Hillsdale, NJ: Lawrence Erlbaum Associates.

Richman, H. B., & Simon, H. A. (1989). Context effects in letter perception: Comparison of two theories. *Psychological Review*, *96,* 417-432.

Richman, H. B., Staszewski, J., & Simon, H. A. (1995). Simulation of expert memory with EPAM IV. *Psychological Review*, *102,* 305-330.

Simon, H. A., & Gilmartin, K. J. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, *5,* 29-46.

Zhang, G., & Simon, H. A. (1985). STM capacity for chinese words and idioms: Chunking and acoustical loop hypotheses. *Memory and Cognition*, *13,* 193-201.