Chapter Four

THE ROLE OF MEMORY IN COGNITION: WORKING MEMORY

The study of STM sprang from concern for a number of practical problems, such as the attempt by Jacobs to measure the mental capacity of his pupils, Broadbent's interest in the division of attention and its implications for jobs such as that of the air traffic controller, and Conrad's concern for the memorability of telephone numbers and postcodes. By the late 1960s however, the study of STM had become very much laboratory-bound; it had produced a plethora of novel laboratory techniques and detailed models and theories, often expressed mathematically. Although based almost entirely on laboratory results, none the less the Atkinson and Shiffrin (1968) model did make claims for the general importance of the short-term store. It assumed that the STS acted as a *working memory*, a system for temporarily holding and manipulating information as part of a wide range of essential cognitive tasks such as learning, reasoning and comprehending. Such a view would probably have been quite widely held during the 1960s, although there was little effort to test it directly. The present chapter describes one attempt to investigate the role of short-term storage in a range of tasks and situations, and to ask whether STS really does serve as a general working memory. In attempting to answer that question, the earlier concept of a unitary STS is challenged and replaced by a related but more complex concept, that of a multi-component working memory model. This attempts to account for both the evidence that fitted the earlier STS model and also those features that were problematic. In addition, the concept of working memory attempts to highlight the role of temporary storage in other cognitive tasks such as reasoning, comprehension and learning.

In their levels of processing framework, Craik and Lockhart (1972) continued to assume that primary memory played an important role in cognition, but did they necessarily need to? Given that their framework would deal with coding effects quite effectively, could handle the absence of long-term learning following maintenance rehearsal, and given the uncertainties surrounding the explanation of recency effects, is there any need to assume a short-term store? Even more pressingly, if patients with a severe deficit in short-term storage are apparently otherwise unimpaired and capable of living a full and rich life, is the study of STM anything other than a cul-de-sac in the short but tortuous history of human experimental psychology?

Graham Hitch and I decided to try and tackle this problem by asking the basic question of "What is STS for?" We decided that if the answer was that it merely served to keep experimental psychologists occupied, we would choose to occupy ourselves in other ways.

TESTING THE WORKING MEMORY HYPOTHESIS

A widely held assumption was that STS acts as a temporary working memory that helps us perform a range of other cognitive tasks (Atkinson & Shiffrin, 1968; Hunter, 1957; Newell & Simon, 1972); the concrete evidence for such a view was remarkably sparse. We decided to test it by using a dual-task technique whereby the subject is required to perform one task that absorbs most of the capacity of his working memory, while at the same time performing each of a range of tasks such as learning, reasoning and comprehending that are assumed to be crucially dependent upon working memory. If the assumption is correct, then performing a concurrent STM task should lead to a dramatic impairment in performance.

We selected digit span as our concurrent memory task; although many different models of STS exist, virtually all of them assume that it has a limited storage capacity, and that this capacity is used in performing the standard immediate serial recall task. We began rather tentatively by requiring our subjects to remember only one or two digits while reasoning or learning, but much to our surprise and that of our subjects, they proved to be only minimally encumbered by a few additional items, and we therefore moved to rather heavier concurrent digit loads of three or six items.

We also began by presenting the digits, requiring the reasoning or memory task, and then asking for digit recall. Under these circumstances, however, we found that subjects tended to adopt a strategy of rapidly rehearsing the digits, then switching attention to the reasoning or learning task, before returning to pick up whatever they could of the trace of the digits. As such, we were not obtaining a measure of *concurrent* processing so much as observing the effect of alternating the two tasks. We therefore opted for a policy of requiring the subject always to continue to rehearse the digits out loud, hence ensuring that they were performing both tasks simultaneously. We assumed that if STS serves as a limited-capacity working memory that is used in reasoning or learning, then loading STS with a concurrent task of remembering digits should impair performance. The larger the number of digits being held, the greater the amount of working memory capacity that should be absorbed, and the greater the interference with reasoning or learning performance.

In one study, subjects were required to remember number sequences ranging from zero to eight digits in length, while at the same time performing a reasoning test. This involved verifying a series of sentences each of which purports to describe the order in which two successive letters, *A* and *B*, were presented. The subject's task is to decide whether the sentence correctly describes the order or not. Examples ranged from simple active declarative sentences such as *A follows B – BA* (true) to more complex sentences involving passives and/or negatives such as *B is not preceded by A – AB* (false).

This particular reasoning test is one based on some of the early developments in psycholinguistics which showed that the more complex the sentence, the longer the decision time. I initially devised it to provide a simple robust reasoning test that could be performed underwater as part of a study on the effects on the mental efficiency on deep-sea divers of nitrogen narcosis, the drunkenness that one experiences on breathing air at depth. It proved to be a valid and reliable correlate of verbal intelligence (Baddeley, 1968) and to be sensitive to nitrogen narcosis (Baddeley & Flemming, 1967).

Figure 4.1 shows the effect of concurrent memory load on the speed and accuracy with which subjects performed the syntactic reasoning test. Two points should be noted; first of all, the reasoning time increases clearly and systematically with concurrent memory load, just as a working memory hypothesis would predict. Secondly, however, note that the effect is far from catastrophic; requiring a subject to concurrently rehearse eight items, which in many cases was more than could accurately be maintained, leads to an increase in latency of only about 35%. Even more strikingly, note that error rate remains constant at around 5%. It is not easy to account for this pattern of results if one assumes that working memory involves a single unitary store whose limited capacity is likely to be totally absorbed when the limit of memory span is reached. On this assumption, a concurrent load of eight digits should cause reasoning performance to break down completely. It clearly does not.

A broadly similar pattern of results was obtained across a range of other cognitive tasks. In one study, the free recall of lists of unrelated words was studied when they were accompanied by a concurrent digit span of zero, three or six items. Performance on the earlier part of the serial position curve, normally associated with long-term learning was impaired, but by no means obliterated by the concurrent load of six digits, while a three-digit load had no significant effect on performance. As mentioned earlier, not even

FIGURE 4.1

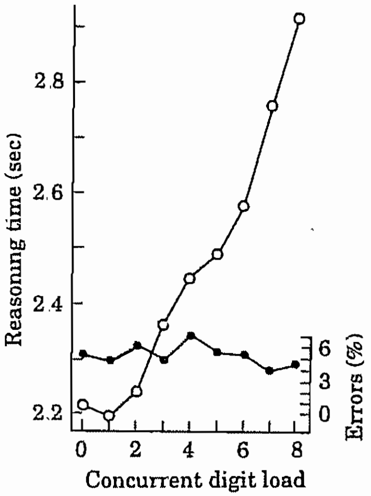

Speed and accuracy of grammatical reasoning as a function of concurrent digit load. From Baddeley (1986)

the six-digit load had any effect on the magnitude of recency, suggesting that recency and digit span may well reflect the operation of different memory systems.

In another study, comprehension of prose passages was studied in subjects who were concurrently remembering sequences of zero, three or six unrelated digits. Level of comprehension was significantly impaired by the six-digit, but not the three-digit load (Baddeley & Hitch, 1974).

Finally and most unexpectedly, despite a clear effect on learning, concurrent digit span during *retrieval* from long-term memory was found to have no effect on accuracy of performance, although it did produce an increase in retrieval latency (Baddeley, Lewis, Eldridge, & Thomson, 1984). The requirement to remember and recite a six-digit number had no effect on the accuracy of recalling or recognizing lists of words, whether tested by free recall or paired-associate learning (see Table 4.1). This somewhat surprising result suggests that whatever system is responsible for holding digits in immediate memory might be, it does not play the crucial role in retrieval assumed by many earlier models of memory including that of Atkinson and Shiffrin (1968) and that of Rumelhart, Lindsay, and Norman (1972).

A WORKING MEMORY MODEL

How well then, did the concept of a working memory survive this initial exploration using the dual-task paradigm? With the exception of retrieval, all the other tasks showed evidence for impaired

FIGURE 4.2

A simplified representation of the working memory model

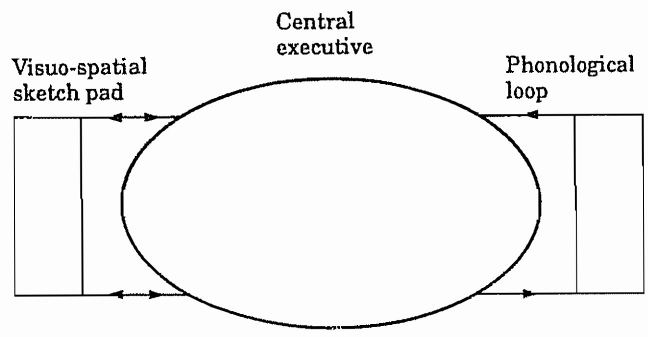

performance, with the degree of impairment increasing with concurrent digit load, a pattern of results that is quite consistent with the working memory hypothesis. On the other hand, the degree of impairment is far from dramatic, particularly bearing in mind that a concurrent load of six digits would approach the span of most of our subjects, and on assumptions such as those of the modal model ought to leave very little processing capacity left for reasoning, learning, comprehending or retrieving. The simplest way out of this paradox seemed to be to abandon the assumption of a unitary STS, and accept that the limits of digit span may be set by one of a number of subsystems, leaving other components of working memory relatively unimpaired.

On the basis of the evidence from these and other tasks, we proposed a model of working memory in which a controlling attentional system supervises and coordinates a number of subsidiary slave systems. We termed the attentional controller the central executive and chose to study two slave systems in more detail, the articulatory or phonological loop which was assumed to be responsible for the manipulation of speech-based information, and the visuo-spatial scratchpad or sketchpad, which was assumed to be responsible for setting up and manipulating visual images. A simple representation of the model is shown in Figure 4.2. I shall begin by discussing the slave systems before going on to talk about the more difficult task of exploring the central executive.

THE PHONOLOGICAL LOOP

We postulated this particular subsystem in order to give an account of the very substantial evidence for the importance of speech coding in STM. It is probably the most extensively worked out component of the model, partly because I suspect it is one of the simpler components, and partly because it is concerned with an area where considerable data already existed. Because of this, it offers a good example of a particular approach to theorizing, namely that of attempting to constrain possible models by using a rich and robust pattern of results, any one of which is capable of being explained in several different ways, but which together place major constraints on possible explanations. The theoretical aim in the short term is to provide a simple account of all the data; such an account need not, and rarely is quantitative and precise, but it represents the basic structure that any more detailed model will need to encompass.

Since the pattern of data is relatively complex, I will begin by giving a brief overview of the assumed structure of the phonological loop system, followed by a description of the individual phenomena, after which an overall mapping of the phenomena onto the model will be suggested.

The phonological loop is assumed to comprise two components, a phonological store that is capable of holding speech-based information and an articulatory control process based on inner speech. Memory traces within the phonological store are assumed to fade and become unretrievable after about one-and-a-half to two seconds. The memory trace can however be refreshed by a process of reading off the trace into the articulatory control process which then feeds it back into the store, the process underlying subvocal rehearsal. The articulatory control process is also capable of taking written material, converting it into a phonological code and registering it in the phonological store. This simple model of a phonological store served by an articulatory control process can give a coherent account of the following phenomena:

Evidence for the Loop

The Phonological Similarity Effect

As we saw earlier immediate serial recall is impaired when items are similar in sound or articulatory characteristics, hence PGTVCD will be harder to remember than RHXKWY (Conrad & Hull, 1964; Baddeley, 1966a). The question of whether the crucial aspect of similarity is at the level of sounds, phonemes or articulatory commands is one that has created a good deal of discussion over the years, without reaching any very satisfactory conclusion, since the various measures are all extremely highly correlated (e.g. Hintzman, 1967; Wickelgren, 1969). Virtually all the terms used in this respect, including "acoustic", "phonemic" and "phonological" can be taken to imply a particular position on this issue. The present use of the term "phonological" is, however, meant to be relatively neutral on the issue of exactly what level of speech coding is involved. The phonological similarity effect is assumed to occur because the store is based on a phonological code, hence similar items will have similar codes. Recall will require discriminating among the memory traces. Similar traces will be harder to discriminate, leading to a lower level of recall.

The Unattended Speech Effect

Colle and Welsh (1976) carried out a study in which subjects attempted to repeat back sequences of visually presented numbers. In one condition, immediate serial recall was accompanied by the sound of someone reading a passage in German, a language the subjects did not understand. Nevertheless, performance on the immediate memory task showed a clear decrement.

A colleague, Pierre Salamé, and I independently stumbled across the same effect a few years later. Pierre had been working on the effects of noise on memory, in France, and on a collaborative working visit to Cambridge decided that it would be interesting to extend the range of possible distractors to include spoken words. He predicted that being meaningful, the words would be particularly distracting, while I suspected that the subject would be quite capable of ignoring them, producing a negative result.

We therefore set up an experiment in which subjects attempted the immediate recall of nine visually presented digits which were presented either in silence, or accompanied by spoken words or spoken nonsense syllables, both of which the subject was instructed to ignore. Pierre predicted that performance would be more severely disrupted by the words than by nonsense; I predicted no disruption from either. We were both wrong. Performance was disrupted to an equal extent by both words and nonsense syllables. We concluded that the unattended material was gaining access to the phonological store, a store that holds phonological but not semantic information.

This conclusion was reinforced by a subsequent experiment in which subjects again attempted to remember visually presented digit sequences, this time against a background either of other digits, or of other words made up from the same phonemes as digits (e.g. *tun*, *woo* instead of *one*, *two*). A third condition involved ignoring words that were phonologically dissimilar disyllables (e.g. *happy*, *tipple*), while a fourth comprised a silent control condition. The disyllables caused some disruption but not so much as the monosyllables having the same phonological characteristics as digits. These did not, however, differ in their degree of disruptiveness from actual digits, again suggesting that the store contains phonemic information but does not represent items at a word level, otherwise the digits would have been expected to be more disruptive than the non-digits made up from the same phonemes.

Can any sound gain access to the phonological store? The evidence suggests not. In one study, for example, we compared the effects of unattended speech with that of unattended noise on immediate serial recall of digits. We found a clear effect of unattended speech, but no effect of noise, even when the noise was pulsed so as to give the same intensity envelope as continuous speech (Salamé & Baddeley, 1987; 1989). In some ways the effect resembles auditory masking, but in others it does not. The effect is for example unaffected by the intensity of the unattended speech, provided that it is clearly audible (Colle, 1980; Salamé & Baddeley, 1987).

What about unattended music? We studied this in an experiment in which the subject again tried to recall sequences of visually presented digits, this time against a background of either vocal or instrumental music. Whether the vocal music came from nineteenth-century opera in an unfamiliar language, or from a current pop star singing in the subject's native language, the disruption was the same and approximately equivalent to that produced by unattended speech. In the case of instrumental music, the effect was present but rather less marked, and again was uninfluenced by whether it was represented by modern or classical pieces (Salamé & Baddeley, 1989).

What are the practical implications of our results? Should one definitely avoid studying with the radio on? At present we have not explored a sufficiently wide range of tasks to come up with firm recommendations, but it seems that the effect rather specifically impairs tasks that heavily involve the phonological store. We have not so far obtained any indication to suggest that reading comprehension (Baddeley, Eldridge, & Lewis, 1981) or free recall learning (Salamé & Baddeley, unpublished) are impaired by meaningless unattended speech. On the other hand, evidence is beginning to appear suggesting that if the material is meaningful, and evokes at least some of the listener's attention, then impairment in comprehension and/or retention of prose will be observed (Martin, Wogalter, & Forlano, 1988).

In the past, experimental investigations into "noise pollution" have all too frequently opted to study meaningless white noise, regarding sound intensity as the main variable of study. Our results suggest that the qualitative nature of the noise may be a rather more important factor when it comes to disrupting working memory. However, the area clearly needs a good deal more investigation using a wider range of tasks and a wider range of potentially disrupting sounds before valid conclusions can be drawn.

The Word-Length Effect

Another powerful determinant of immediate memory span is the spoken duration of the words presented. Hence most subjects would relatively easily remember a sequence of five monosyllabic words such as *wit*, *sum*, *harm*, *bag*, *top*, but would have considerable difficulty in repeating back a sequence of polysyllables such as *university*, *opportunity*, *aluminium*, *constitutional*, *auditorium*. Figure 4.3 shows the relationship between word length, reading rate and memory for words ranging in length from one to five syllables. Figure 4.4 shows the relationship between total spoken duration and probability of recall. The results fall on a straight line which can be reinterpreted as indicating that memory span represents the number of items of whatever length that can be uttered in about two seconds. As one might expect from this, there is a correlation between the rate at which a subject speaks and his or her memory span (Baddeley, Thomson, & Buchanan, 1975).

Is the crucial feature spoken duration or number of syllables? Duration appears to be the critical variable since sequences of words that tend to have long vowels and be spoken slowly such as *Friday* and *harpoon* lead to somewhat shorter spans than words

FIGURE 4.3

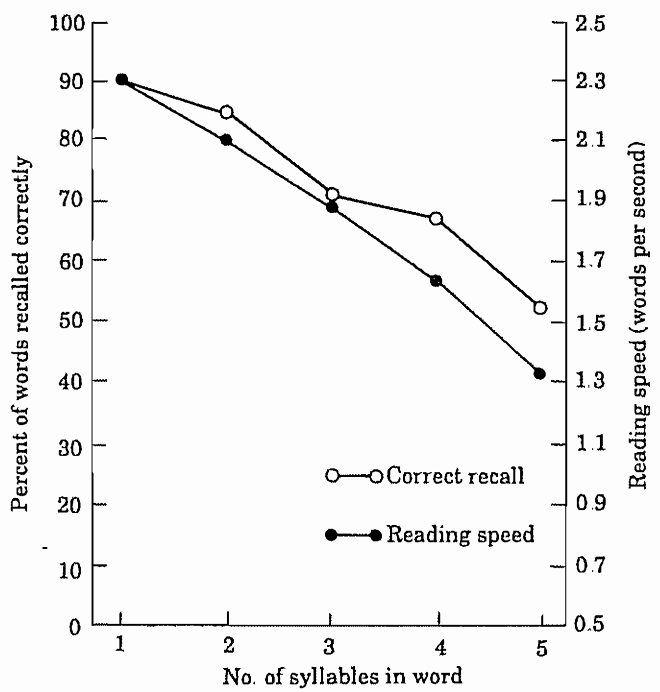

Relationship between word length, reading rate and recall. Long words take longer to rehearse and produce lower memory spans. From Baddeley, Thomson, and Buchanan (1975b).

FIGURE 4.4

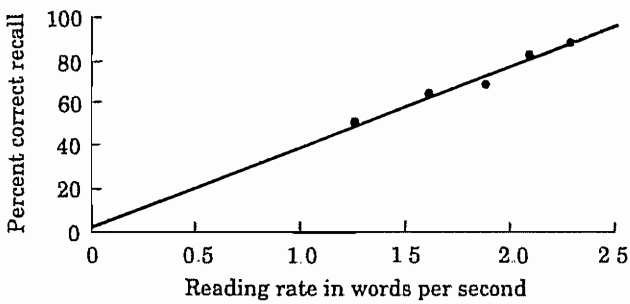

The relationship between reading rate and recall observed by Baddeley et al (1975b)

TABLE 4.1

Set of Words Used by Baddeley, Thomson, and Buchanan (1975)

Mumps	Measles	Leprosy	Diphtheria	Tuberculosis
Stoat	Puma	Gorilla	Rhinocerous	Hippopotamus
Greece	Peru	Mexico	Australia	Yugoslavia
Maine	Utah	Wyoming	Alabama	Louisiana
Zinc	Carbon	Calcium	Uranium	Aluminium

with the same number of syllables and phonemes that can be spoken more rapidly (e.g. *wicket*, *bishop*).

Ellis and Hennelly (1980) used this effect to interpret an anomaly in the detailed intelligence test results for Welsh-speaking children. Such children appear to have a consistently lower performance on the digit span subtest of the Wechsler Intelligence Scale than their English-speaking contemporaries. Could this indicate some strange genetic quirk of the Welsh to offset perhaps their prowess at choral singing and rugby playing? Ellis and Hennelly proposed a more prosaic interpretation, observing that the digit names in Welsh, although having the same number of syllables as in English, tend to have longer vowel sounds and take longer to say. They tested this using bilingual Welsh- and English-speaking subjects. Their subjects proved to have a poorer span in their native language of Welsh than in English, but as predicted, their spans were equal when measured in terms of spoken time.

When the subjects were prevented from rehearsing by the requirement to utter an irrelevant sound, the difference between span in the two languages disappeared. Was the effect purely limited to the memory span? Apparently not since their subjects also showed some signs of slower performance and higher error rate in mental arithmetic using Welsh digits. Later research by Naveh-Benjamin and Ayres (1986) has extended the work of Ellis and Hennelly across a range of different languages, and as Figure 4.5 shows obtaining a clear relationship between memory span, and the time it takes to articulate the digits one to ten in that language.

However, the record so far for speed of articulation goes to Chinese for which Hoosain and Salili (1988) report a mean articulation rate of 265 milliseconds per digit compared to Ellis and Hennelly's report of 321 milliseconds for English and 385 milliseconds for Welsh. Mean digit span was no fewer than 9.9 for Chinese subjects compared to a mean of 6.6 for English, and 5.8 for Welsh. Hoosain and Salili also report a correlation between memory span and mathematics exam grades of 0.38. They further report that recitation of multiplication tables is much faster in the case of Chinese undergraduates using Cantonese (mean time = 64.3 seconds) than for U.S. undergraduates using English (134.2 seconds).

FIGURE 4.5

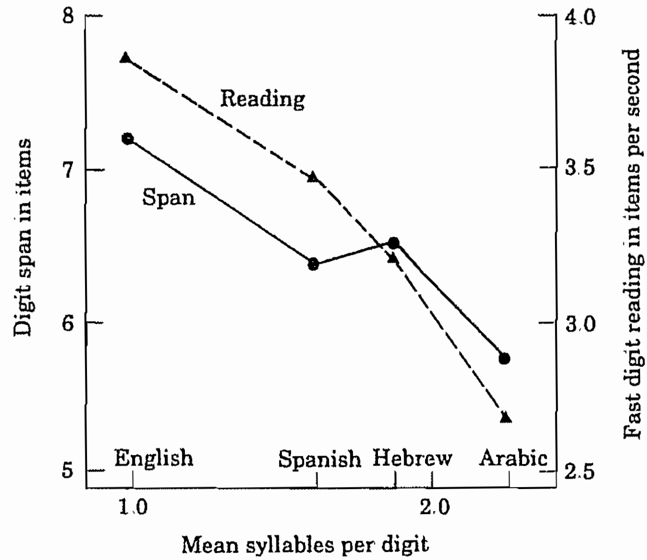

Memory span and reading rate for digits in four different languages. From Naveh-Benjamin and Ayres (1986). Copyright (1986) The Experimental Psychology Society

Finally, Hoosain and Salili report a study by Chan (1981) who compared two groups of students in Hong Kong, both from the same Chinese-speaking primary school. One group elected for education in English, the other in Chinese. The groups were equivalent in performance when they transferred schools at age 12. Five years later, those who went to Chinese-speaking schools tended to have poorer grades in everything other than Chinese and mathematics. Such a result is open to a range of possible interpretations, but does support the possibility of a link between language, the phonological loop and mathematics, reinforcing other evidence indicating an involvement of the phonological loop in counting (Logie & Baddeley, 1987) and mental arithmetic (Hitch, 1978).

One of the clearest and most reliable features of the development of memory in children is the tendency for digit span to increase systematically with age. Nicolson (1981) made the interesting suggestion that this might be due to a tendency for older children to rehearse faster. He studied the speed at which children of different ages could articulate and plotted their memory span as a function of this, finding a very clear relationship.

This finding has subsequently been replicated and further extended by Hulme, Thomson, Muir, and Lawrence (1984) and by Hitch, Halliday, and Littler (1984) in a series of studies in which children of various ages were tested for immediate serial recall of

FIGURE 4.6

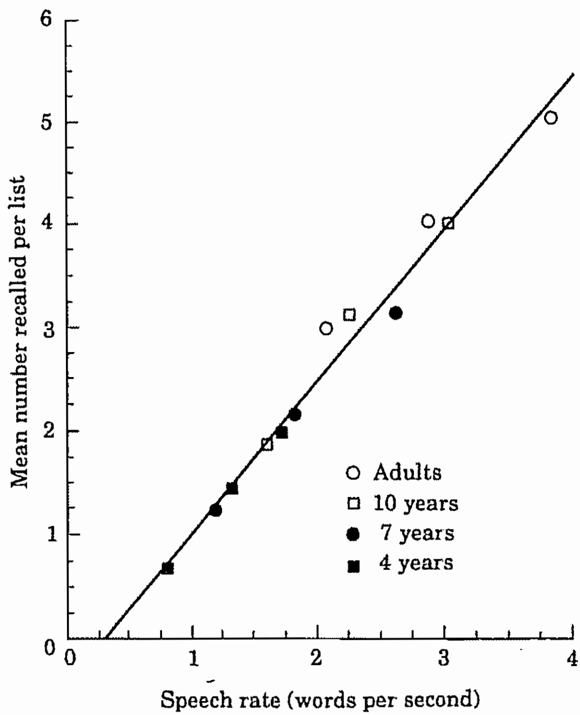

The relationship between word length, speech rate, and memory span as a function of age From Hulme. Thomson, Muir, and Lawrence (1984)

items that had names varying in length. When presentation was auditory, length had an effect down to children as young as four. As Figure 4.6 shows, when memory span is plotted against articulation rate, the data from all the item lengths and ages fall on the same straight line, suggesting that in this task at least, increased age enhances performance simply because subjects articulate more rapidly.

The presence of a word-length effect implies some form of subvocal rehearsal, and as such appears to conflict with other data from a study for example of free recall in children which seems to suggest that cumulative rote rehearsal is not developed until a somewhat later age (Ornstein, Naus, & Stone, 1977). A subsequent experiment suggests that rehearsal may be very dependent upon a particular experimental situation. For example, if items are presented visually as drawings, rather than spoken, then a word-length effect does not appear in children of the age of 6 and 8, although it is present by the age of 10 (Hitch et al., 1984). It seems likely that the tendency to repeat back either overtly or covertly, items that are heard, is a rather early development, possibly playing an important role in acquiring spoken language. It appears that converting a visual item into a phonological code, and rehearsing that code is something that is acquired considerably later.

How should the word-length effect be interpreted? The simplest account might be to suggest that the process of overt or covert articulation involves setting up and running speech motor programs which operate in real time, with the result that the longer the word the longer it takes to run off. If we assume that this process of subvocal rehearsal has the function of maintaining items in the phonological store by refreshing their fading traces, then the faster it can run, the more items will be maintained and the longer the memory span. If we assume that the memory fades, then the memory span will be determined by the number of items that can be refreshed before they fade away. That number, of course, will depend both on how rapidly the trace fades and on how long it takes to articulate each item and hence refresh each memory trace. Data from studies using English, Welsh, Hebrew, Spanish, Arabic and Chinese all give results suggesting that trace decay time is approximately two seconds, although as mentioned earlier, rehearsal time, and consequently span vary widely from one language to another (Hoosain & Salili, 1988).

Articulatory Suppression

Although overt articulation is not necessary for the operation of inner speech, the operation of the phonological loop is disturbed if overt or covert articulation of an irrelevant item is required. Hence, if a subject in a standard digit span task is required to utter a stream of irrelevant sound, such as repeatedly saying the word *the*, span is likely to be substantially lower, whether presentation is auditory or visual. This is assumed to occur because the articulation of an irrelevant item dominates the articulatory control process, hence preventing it from being used either to maintain material already in the phonological store, or convert visual material into a phonological code. It might furthermore have the additional drawback of creating an unattended speech effect by feeding the irrelevant spoken material into the phonological store.

It is important in interpreting the effects of articulatory suppression to consider another possible interpretation, that suppression impairs performance simply because it demands attention (e.g. Parkin, 1988). There are three arguments against this view:

1. Non-articulatory secondary tasks which might reasonably be regarded as similar in level of demand, such as tapping at the same rate as suppression, typically have little or no effect on STM performance (e.g. Baddeley, Lewis, & Vallar, 1984b).
2. Patient P.V., who does not appear to use the articulatory loop in STM tasks, is not impaired in memory performance by suppression (Vallar & Baddeley, 1984a).
3. The complex pattern of results to be described indicates that

suppression does not have a major general effect, but rather specifically affects phonological and articulatory coding (Baddeley et al., 1984b). It is, none the less, always wise in studies of articulatory suppression to include a tapping condition to control for any general attentional effects on performance.

The effect of articulatory suppression is fortunately very robust, and does not appear to be crucially dependent on the items uttered, with different laboratories tending to favor suppression based on different utterances, ranging from *bla bla bla* through *double double double* to over-learned sequences such as counting or reciting fragments of the alphabet. Developmental psychologists tend to use words they think will appeal to children such as *teddy bear teddy bear*, while one U.S. investigator uses *cola cola cola* which suggests some interesting advertising possibilities within this paradigm.

If we assume that articulatory suppression cuts out the process of subvocal rehearsal, then it ought to interact in predictable ways with some of the other variables we have described, as indeed it does. In the case of phonological similarity, articulatory suppression removes the effect when material is presented visually, presumably since it prevents the visual code from being converted by subvocalization into a phonological code that can be registered in the store. With auditory presentation, however, the phonological similarity effect remains, presumably because the spoken material has direct access to the phonological store without need of the articulatory control process (Baddeley et al. 1984b).

As with phonological similarity, the unattended speech effect should be disrupted by articulatory suppression, given that the material to be recalled is presented visually. If suppression prevents the subject from subvocally registering the material to be remembered in the phonological store, then memory will be based on some non-phonological store. Corruption of the phonological store by unattended speech should hence not affect performance. This is indeed what was found (Salamé & Baddeley, 1982). With auditory presentation however, recall of the items to be remembered will depend on the phonological store, and unattended speech should therefore impair performance, which it does (Hanley & Broadbent, 1987).

What effect should articulatory suppression have on the word-length effect? Since the effect depends directly on subvocal articulation, then articulatory suppression should abolish it, regardless of whether presentation is visual or auditory. If subjects are prevented from rehearsal, it should not matter whether the material is fast or slow to rehearse.

Initial results on this point seemed rather worrying, since Baddeley, Thomson, and Buchanan (1975b) found that articulatory suppression disrupted the word-length effect with visual, but not with auditory presentation. This result went against the predictions of the model which assumes word length to influence rehearsal rate but not storage. Further experimentation however, revealed a critical flaw in the initial experiment using auditory presentation. Articulatory suppression had occurred during presentation but not during recall. It appears that subjects were rapidly rehearsing the auditorily presented items before and during subsequent recall. When suppression is required during both presentation and recall, no significant word-length effect is found (Baddeley et al., 1984b).

A Summary of the Evidence

As explained earlier, the mode of theorizing in working memory has involved taking a relatively complex pattern of data and attempting to fit it into as simple a conceptual structure as possible. In the case of the articulatory loop it has involved a phonological store which will hold information for about two seconds, together with an articulatory control process. This process refreshes items in the store by means of subvocal rehearsal; it is also capable of subvocally recoding printed material, hence registering it in the phonological store.

The phonological similarity effect occurs because the store is based on phonological coding. Similar items have easily confusable codes, leading to impaired performance. Articulatory suppression prevents visual material being recoded, but has no effect on the coding of auditory material, which hence continues to show a similarity effect.

The unattended speech effect is assumed to occur because spoken material gains obligatory access to the phonological store, which is corrupted by the presence of irrelevant material. Suppression prevents the unattended speech effect occurring with visually presented material, since it stops such material being fed into the phonological store. Whether the store is or is not corrupted therefore becomes irrelevant.

Finally, the word-length effect is removed by articulatory suppression, whether material is presented auditorily or visually. Since the word-length effect is dependent on the operation of the articulatory control process, when this system is pre-empted by suppression, word length ceases to be an important variable.

Chunking and the Phonological Loop

The essence of the phonological loop hypothesis is that memory span will depend on rate of rehearsal, being approximately equivalent to the number of items that can be spoken in two seconds. Hence number of items recalled will be a function of how long they take to articulate. Where then does this leave Miller's magic number seven, which suggests that memory span will reflect a constant number of chunks, regardless of the characteristics of those chunks? This question was addressed directly by Herbert Simon who in addition to a distinguished range of other activities in cognitive science had adopted and developed the chunking hypothesis as an important feature of human cognition (Simon, 1974). He and a group of Chinese colleagues took advantage of some of the features of the Chinese language to explore the articulatory loop and chunking hypotheses in more detail (Zhang & Simon, 1985; Yu, Zhang, Jing, Peng, Zhang, & Simon 1985).

In one study, Zhang and Simon directly pitted the two hypotheses against each other, using three types of material which were equivalent in each comprising familiar chunks, but differed in ease and speed of pronunciation. One set of material comprised radicals, the complex components which go to make Chinese characters and words. There are about 200 radicals in the Chinese language, and they are likely to be highly familiar to their Chinese subjects since they are, for example, used for indexing dictionaries. They do not, however, have commonly used oral names. The second set of material used comprised Chinese characters, each of which was made up from two radicals, with each having definite single-syllable pronunciation. The third set of material comprised Chinese words, each comprising two characters and having two syllables in their pronunciation. The different types of material are shown in Figure 4.7.

A simple chunking hypothesis would predict no difference between the three sets of material, since in each case the constituent items comprised familiar chunks. The phonological loop hypothesis on the other hand would predict very poor performance for the radicals which have no familiar name, with somewhat better performance for the disyllabic words, and best performance for the monosyllabic characters. As Table 4.2 shows, it is exactly what was observed. Further evidence for the phonological loop interpretation came from intrusion errors, of which almost half were homophones, items that have the correct pronunciation but are written differently.

FIGURE 4.7

Chinese radicals, characters and words used by Zhang and Simon (1985) Reprinted by permission of the Psychonomics Society Inc

Set	Radical	Character	Word
1	爫	爱	爱人
2	𠂇	友	友谊

TABLE 4.2

Mean STM Span for Three Types of Chinese Symbol Sequences (Data from Zhang & Simon, 1985)

Type of Item	*Mean*	*s.d*
Radicals	2.71	0.52
Characters	6.38	1.08
Words	3.83	0.75

Chinese has a very large number of homophones, and since the pictographic script is not based on the sound of the items depicted, such homophones are typically written quite differently. Zhang and Simon took advantage of this in order to explore further the role of phonological coding in a second study. This used the material shown in Figure 4.8 which comprises nine characters, all of which are pronounced *gong* with high tone in Chinese. Memory span for

FIGURE 4.8

Set of homophonic characters, all pronounced "gong", with high tone, but all with different appearance and meaning, used by Zhang and Simon (1985) to study the role of phonological similarity in immediate memory. Reprinted by permission of the Psychonomics Society Inc

Chinese character	English translation
工	work, labour
弓	bow
公	public, common
功	meritorious service
攻	attack, accuse
供	supply
宫	palace, temple
恭	respectful
龚	a surname

these items was tested together with memory span for the nonpronouncable radicals described previously. Span for the radicals was 3.00 and for the homophones 2.83, in contrast to the span for characters in the previous study of 6.38. This is of course a very low level of performance, and indeed one subject complained "It seems that my memory doesn't work today". Presumably the items recalled were being remembered on the basis of either a visual or semantic code.

In a third experiment, Zhang and Simon explored memory span for items varying in number of syllables, testing memory for characters, comprising one syllable, words of two syllables and idioms comprising four syllables. Table 4.3 shows the results they obtained measured in terms both of mean number of chunks recalled and mean number of syllables. It is clear that while span does not represent a constant number of chunks, it does not either represent a constant number of syllables. Syllables within chunks tend to lead to faster articulation than do syllables that comprise separate chunks.

Zhang and Simon therefore propose that span is determined by rehearsal rate, but that this in turn depends on three factors, the interval of time (a milliseconds) required to bring each chunk into the articulatory mechanism, an interval of time (b milliseconds) required to articulate each syllable in the chunk beyond the first, and S the average size of a chunk in syllables. This yields an equation that can be used to express either T, the duration of the underlying storage parameter, or C, the STM capacity measured in chunks. They are

$$T = C[a+b(S-1)], \quad \text{or} \qquad (1)$$
$$C = T/[a+b(S-1)]. \qquad (2)$$

Zhang and Simon show that these equations fit a wide range of experimental results collected in Chinese, and also the data on word length and memory in English reported by Baddeley, Thomson, and Buchanan (1975).

In general, these experiments, taking advantage of some of the

TABLE 4.3

Mean STM Span for Three Kinds of Chinese Symbol Sequences. (Data from Zhang & Simon, 1985)

Type of Item or Chunk	*No. of Syllables per Chunk*	*Mean Recall*	
		Chunks	*Syllables*
Characters	1	6.58	6.58
Words	2	4.58	9.16
Idioms	4	3.00	12.00

intriguing characteristics of the Chinese language, produce results that support the phonological loop hypothesis. At the same time they suggest ways in which the more general concept of chunking in STM might be incorporated into the model.

Patients with Impaired STM

A bonus from the phonological loop model is that it offers a straightforward explanation of the memory deficit shown by STM patients. If one assumes a deficit in the phonological store, then this is able to explain both their impaired memory span and their comparatively normal cognitive performance on other tasks such as long-term verbal learning, where one might expect semantic coding to be more important than phonological. One Italian patient, P.V., with a very pure and specific deficit in auditory STM performance was tested with a view to exploring the extent to which her deficit could be explained within the working memory framework (Vallar & Baddeley, 1984a). With the exception of her STM deficit, P.V. appeared to be intellectually entirely normal, with a high level of verbal and performance I.Q., excellent long-term memory, and no apparent problems of speech or language (Basso et al. 1982). P.V.'s immediate memory was influenced by phonological similarity with auditory, but not with visual presentation. She showed no evidence of a word-length effect, and no effect of articulatory suppression. Her speech output appeared to be quite normal as measured both by studying the distribution of pauses in normal speech, and as measured by her capacity to recite the alphabet or count as rapidly as normal control patients of equivalent age and background.

We interpreted her deficit as an impairment though not complete disruption of the phonological storage component of the articulatory loop. We assumed that disruption was not complete since her performance with auditory presentation, though impaired was not completely disrupted, and did show clear evidence of a phonological similarity effect. We assume, however, that she does not attempt to feed visually presented material into the phonological store by the process of articulation. The evidence for this assumption comes from the absence of a phonological similarity effect with visual presentation, the lack of a word-length effect and the absence of any effect of articulatory suppression on performance. We assume that her STM deficit does not stem from an articulatory deficit since her capacity for overt articulation appears to be normal, and since she appeared to be able to make phonological judgments about printed words, deciding for example whether the names of two pictured objects rhymed or not, or where the stress on a particular printed word occurred. We assume that she does not use the articulatory rehearsal process simply because it would feed information into a grossly defective store which would do little to enhance performance.

Dysarthria and the Nature of Inner Speech

To what extent does this process of rehearsal need to involve the overt activity of the speech musculature, and to what extent can it be maintained at some higher more programmatic level? We certainly do not need to rehearse out loud, although it is possible that silent rehearsal still involves some subvocal activity that can perhaps be detected by electromyography, a process whereby electrical activity in the underlying speech musculature can be monitored. A number of studies have attempted to explore the role of subvocal speech in reading using this approach (e.g. Hardyk & Petrinovich, 1970), although the interpretation of the results of such studies remains open to question. Failure to detect any activity could for instance simply reflect insufficient sensitivity in the equipment. Furthermore, if effects are detected, they could reflect a general overflow of activation, rather than an essential feature of subvocal rehearsal. Consequently, electromyography has not featured prominently in recent discussions on the role of inner speech in memory.

An alternative is to examine the memory performance of subjects who are *dysarthric*, that is patients who have lost the capacity to control their articulatory muscles as a result of brain damage. Dysarthria typically results from damage to the brain stem or peripheral aspects of speech control, and as such should be distinguished from *dyspraxia*, problems at the level of setting up and running the motor programs necessary for speech, and *dysphasia*, which would typically involve more central disruption of the capacity to produce and/or comprehend language.

A colleague, Barbara Wilson, and I were able to study the memory performance of a group of dysarthric patients, with relatively severe but peripheral disruption to their capacity to generate speech, and one anarthric patient who was totally unable to make any sound other than an inspiratory groan. This latter patient completely lacked the capacity to articulate, and yet his language capabilities were unimpaired, as indicated both by his comprehension performance, and by his language production using a simple keyboard device.

We tested this patient's memory and found first of all that he had a comparatively normal digit span of six items. Furthermore, he showed a very clear phonological similarity effect, for both visual and auditory material, suggesting that he was using the phonological store in a normal way. We tested for the presence of subvocal rehearsal by means of the word-length effect, and found this to be quite normal. Finally, we assessed his capacity for making phonological judgments on printed material, requiring him to decide whether two items would sound the same if spoken. We tested both words (e.g. *key-quay*) and non-words (e.g. *frelame-phrelaim*), and also asked the patient to judge whether non-words were homophonous with real words (e.g. *oshun*). He was able to perform all of these tasks accurately and with no apparent difficulty. In short then, our dysarthric patient appears to have normal inner speech. Broadly similar results have subsequently been obtained by a range of other studies and are reviewed by Logie, Cubelli, Della Sala, Alberoni, and Nichelli (in press).

We interpreted our results as suggesting that the articulatory control process does not depend upon peripheral speech musculature for its operation. Presumably some form of motor program can be run at a central level, despite the absence of peripheral feedback. Our subjects had all previously had normal language and speech, raising the further question of whether the feedback from overt speech is necessary for a child to learn to use subvocal rehearsal and the articulatory loop. Some recent research by Biship and Robson (1989) suggests that it is not. They studied the memory performance of children who had been anarthric since birth, and who had never in their lives been able to articulate speech. Somewhat surprisingly, these children appeared to have normal functioning of the articulatory loop, with relatively normal memory spans and clear evidence of the effects of phonological similarity and of word-length.

It appears then that inner speech is not dependent on outer speech for either its development or its operation. This suggests that the term "phonological loop" is perhaps preferable to "articulatory loop", since the latter seems to imply a direct involvement of articulation.

The fact that inner speech develops under the apparently inhospitable conditions of congential anarthria is intriguing, and suggests that it might perhaps play a rather important role in the development of cognition. The question of what this might be will be discussed next.

WHAT USE IS THE PHONOLOGICAL LOOP?

We have just described in some detail one hypothetical component of one aspect of memory. The simple model presented may be able to give a reasonably economical account of a relatively wide range of laboratory data, but it leaves open the rather crucial question of what function if any is served by this system. Is the articulatory loop anything more than a way of linking together a number of laboratory phenomena? Is it, to use my colleague Jim Reason's blunt but colorful phrase, anything more than "a pimple on the face of cognition"?

Learning to Read

I believe it is, for a number of reasons. First of all, the evidence seems to suggest that the articulatory loop, or some similar system plays an important role in learning to read (Jorm, 1983). If you select a group of children who have a specific problem in learning to read, despite normal intelligence and supportive background, one of the most striking features they have in common is an impaired memory span (Miles & Ellis, 1981). They also, however, tend to perform rather poorly on tasks that do not directly test memory. Such tasks, typically involve phonological manipulation, or require phonological awareness; examples include judging whether words rhyme, or taking a word and deleting the first phoneme before repeating it (e.g. when the subjects hear *spin* they must respond *pin*). Consequently there is some controversy as to whether the deficit underlying the normal development of reading is one of memory, phonological awareness or some third common underlying factor (Bradley & Bryant, 1983; Morais, Allegria, & Content, 1987).

There is, furthermore, clear evidence for a reciprocal relationship between these factors and learning to read, such that learning to read enhances performance on memory span and phonological awareness, which in turn are associated with improvements in reading (Ellis, 1988). Adults who are illiterate as a result of lack of opportunity tend to show impaired phonological awareness, and to improve as they learn to read (Morais et al., 1987). Which comes first then, phonological memory, phonological awareness, or reading?

In the normal development of reading, there is little doubt that these factors interact, but it seems likely that in the case of a minority of children at least, initial reading is handicapped by some form of phonological deficit, a deficit that can be detected before the child has begun to learn to read (Mann & Liberman, 1984). It seems likely that this deficit is related to the development of the phonological loop system, although at present we know too little about it to draw any firm conclusions.

Language Comprehension

Suppose, however, that we do concede that the articulatory loop is useful in learning to read, from an evolutionary viewpoint, reading surely developed too recently for this to offer a plausible explanation as to why an articulatory loop system should have evolved. A more plausible explanation might be to suggest that the phonological loop developed in the process of the evolution of speech production and comprehension. What is the evidence that the phonological store plays a role in speech comprehension? Surely, if this were the case then our STM patients should be unable to understand normal conversation, and hence should be much more handicapped than they in fact are. This depends crucially on exactly what the role of the phonological loop system is in speech comprehension.

One hypothesis proposed by Clark and Clark (1977) suggests that sentence comprehension demands that the whole of each sentence should be held in some temporary store while it is processed grammatically. This hypothesis implies that a whole sentence must be stored before it can be understood. Since our

Italian patient P.V. can understand sentences that are far longer than she can remember, we can reject this view (Vallar & Baddeley, 1984b).

In stark contrast to the extreme dependence of comprehension on STM suggested by Clark and Clark (1977) is the claim by Butterworth, Campbell, and Howard (1986) that language comprehension is quite independent of STM capacity. They studied the performance of a student who had apparently always had a reduced memory span (four digits). They noted that the reading performance of this particular student was highly atypical, resembling the pattern normally categorized as phonological dyslexia. She could read words relatively normally, but had great difficulty in reading even simple non-words such as *JEX* or *FRIMBLE*. It appeared that this subject had indeed had difficulty in learning to read, and appeared to have learnt almost exclusively by a look-say method, allowing her to identify words that had been encountered and learned before, but giving her no phonological skills for pronouncing new and unfamiliar letter sequences (Campbell & Butterworth, 1985). As such, this case gives further support for the idea that normal STS is necessary for learning to read in the usual way, but gives little comfort for the view that auditory *comprehension* is dependent on normal span, since she appeared to have no difficulty in coping with life as a student on a relatively demanding course.

Interpreting these findings is, however, far from simple. The observed span of four digits was low but not as low as most STM patients. Moreover, there is some evidence to suggest that such cases of developmental phonological processing impairment may show unusual neuro-anatomical features that suggest that generalization to the general population may be unwise. Thirdly, there was evidence to suggest that this student was using somewhat different language comprehension strategies from other subjects. For example, comprehension was not impaired by articulatory suppression, making her performance under these conditions better than that of normal students. These issues are discussed further by Howard and Butterworth (in press) and by Vallar and Baddeley (in press).

Giuseppe Vallar and I studied the comprehension of spoken and written discourse by our STM patient, P.V. Although her digit span is only two, she has a sentence span of about six words, and is able to understand simple sentences such as *"Slippers are sold in pairs"* or *"Archbishops are made in factories"*, and answer them rapidly and accurately. Even when the sentences are made longer by adding verbiage, such as *"It is commonly believed and with justification that slippers belong to the category of objects that are bought in pairs"*, she has no difficulty. She does, however, have problems with long sentences where retention and appropriate processing of word order is essential for their correct comprehension. For example, *"The world divides the equator into two halves, the northern and the southern"* or *"It is* *fortunate that most rivers are able to be crossed by bridges that are strong enough for cars"*. On these, she performs at chance, whether she is reading them or hearing them. However, when they are reduced in length (*"The world divides the equator"*; *"Rivers are crossed by bridges"*), performance returns to normal (Vallar & Baddeley, 1984b).

Subsequent more extensive investigation of P.V.'s comprehension reinforces the conclusion that she is able to comprehend normally, provided the material does not require the verbatim retention of information over many intervening words; when it does, her performance deteriorates (Vallar & Baddeley, 1987).

These results suggest that the phonological input store does play a role in comprehension, but possibly only for particularly complex or demanding material. One interpretation therefore, is that it merely acts as a supplementary back-up that plays a secondary role in comprehension, but is not of primary importance. A second possibility, however, is that P.V. is able to comprehend most incoming material reasonably adequately because her phonological store, though impaired, is by no means completely defunct. It may well be that an incoming sentence span of six words is enough to cope with most normal language, whereas a subject with no phonological store whatsoever might conceivably be totally incapable of comprehending speech. One way of exploring these possibilities is to look for a patient with an even lower sentence span than P.V. We were fortunate in locating such a patient, and obtaining his cooperation (Baddeley & Wilson, 1988a).

The patient in question, T.B., was a professional mathematician who reported memory problems following an attack of epilepsy. He was intellectually still functioning at a very high level, but had some long-term memory problems coupled with severe impairment in immediate memory performance, with a digit span of one to two items and a sentence span of three words. T.B. reported that he certainly did have problems in comprehension, saying that typically he could understand the beginning of a conversation, but after the first phrase or so, everything became jumbled. His performance on comprehension tests confirmed this. When given the simple sentences such as *"Bishops are made in factories"*, he could confirm or reject them rapidly and accurately. Once the sentences were made verbose by adding additional phrases, however, he had so much difficulty that he became distressed and we had to stop the test.

Using a series of shorter sentences we found the relationship between sentence length and probability of comprehension shown in Figure 4.9. However, as sentences become longer they also tend to become syntactically more complex; could grammatical complexity be the crucial variable?

We first tried to test this by using visual presentation, on the assumption that the printed word might help substitute for his failing memory. This appeared to be the case since his performance substantially improved, although in the case of long and complex

FIGURE 4.9

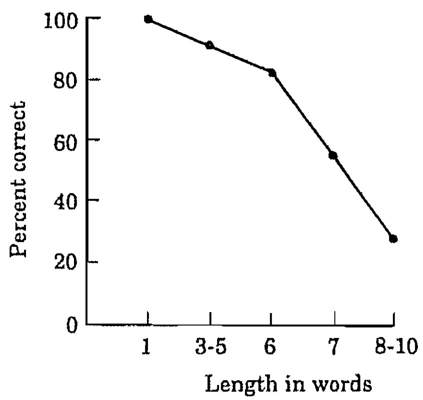

Probability of sentence comprehension by T.B , a patient with imparied STM, as a function of sentence length The data are based on Bishops's TROG (test for reception of grammar) From Baddeley and Wilson (1988a).

sentences, this was at the expense of extremely slow performance. He would hunt to and fro through the sentence as if trying to perform some kind of verbal jigsaw puzzle, before eventually coming up with a response which was more likely to be correct than in the auditory condition, but was still far from perfect.

In a second test of grammatical complexity hypothesis, we began by selecting 24 six-word sentences which he was capable of verifying with almost perfect accuracy. We then added redundant verbiage in the form of adverbs and adjectives. For example the sentence *The boys pick the apples* would be changed to *The two boys pick the green apples from the tree*. Under these conditions, performance dropped to chance. Since the additional verbiage all involved syntactic and semantic constructs that we knew he could comprehend, it seems likely that the crucial factor that led to the drop in performance was the increased memory load.

We have spoken as though the sole factor of importance may have been sentence length. This is almost certainly not the case, since some sentences of a given length were more difficult than others. In general, these appeared to be sentences where memory load was greatest, a good example being self-embedded sentences of the kind *The boy the dog chases is big*. In general, the relationship between syntactic factors, semantic factors and imposed memory load is complex, and it has long proved difficult to separate the influence of these factors empirically (see Baddeley, 1976; Chapter 12). Our data from T.B., however, together with data from other STM patients (Vallar and Shallice, in press) does support the view that the phonological store plays a clear role in comprehension, although a comprehension deficit may only become obvious in some patients when tested with materials placing a particularly heavy load on phonological storage. A patient such as

T.B. whose sentence span is limited to three words however is likely to suffer from substantial problems in comprehending language.

Long-term Phonological Learning

What other role might the articulatory loop play in normal comprehension? One hint was given by the observation that children with developmental dyslexia also tend to have impaired vocabulary and to show difficulty in rote learning such as is involved in acquiring the multiplication tables, or the order of months in the year (Miles & Ellis, 1981). We decided therefore that it might be

FIGURE 4.10

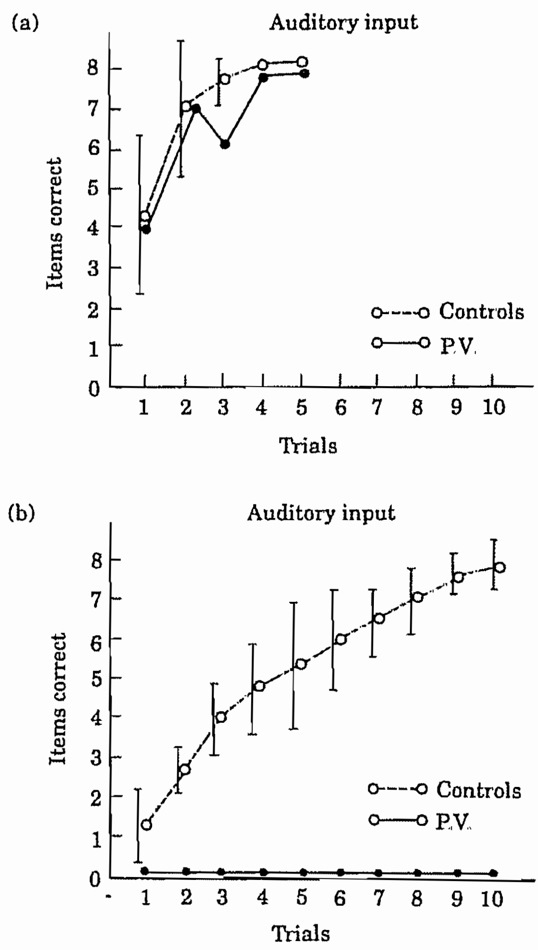

(a) Paired associate learning of auditorily presented word-word pairs by P.V. and by matched control subjects. (b) Performance on word-nonword paired associate learning with auditory presentation; data from P.V. and matched control subjects. From Baddeley, Papagno, and Vallar (1988)

interesting to explore the capacity for learning novel phonological material in our patient P.V.

We knew that P.V. was good at learning lists of meaningful words, but knew nothing of her ability to learn unfamiliar words such as the vocabulary of a foreign language, for example. We therefore decided to try to teach her both pairs of familiar Italian words (her native language), and items of Italian-Russian vocabulary (Baddeley, Papagno, & Vallar, 1988). We found that with auditory presentation, her capacity to learn pairs of Italian words was well within the range of performance shown by controls of equivalent age and intelligence. Her performance on the auditory learning of Italian-Russian pairs on the other hand is shown in Figure 4.10a,b, from which it is clear that she completely fails to learn. When presentation was visual, performance improved somewhat, but was still well below that of control subjects. It appears then that a second function of the phonological loop system is in new phonological learning. If so, then perhaps it plays a crucial role, not only in second language learning, but also in a child's acquisition of his or her native tongue. We were fortunate in being able to explore this possibility in a study of language and memory development in young children that had already begun (Gathercole & Baddeley, 1989).

Acquiring a Vocabulary

A couple of years earlier, Susan Gathercole and I had begun to study a group of children who were classified as "language disordered". They had normal nonverbal intelligence, but delayed development of language skills, and we were interested in particular in their working memory performance. We found that although they were about two years behind their expected performance on reading, vocabulary and spelling, they were four years behind on one task, the simple repetition of nonwords varying in length and complexity. We interpreted this task as one making particularly heavy demands on the phonological loop, suggesting that a phonological loop deficit might be at the root of their other language problems (Gathercole and Baddeley, 1989).

We decided to investigate this task further, and in particular to see if it would allow us to predict which of a sample of 4–5 year-old normal children entering school and about to learn to read, would subsequently prove to have reading difficulties.

We therefore tested over 100 children before they had acquired any reading skills, planning to re-test them at yearly intervals. In addition to our nonword repetition task and measures of reading, we also tested nonverbal intelligence and vocabulary, using a test in which a word is spoken, and the child points to the appropriate picture. At the time of writing, the children have been tested twice. Not enough children have yet learned to read to allow the original question to be answered, but we were able to follow up our hypothesis about the role of the articulatory loop in vocabulary

TABLE 4.4

Correlations between Vocabulary Scores at Age 4 and other Variables

Measures	*Correlation Coefficient*	*Simple Regression (% Variance)*	*Stepwise Regression (% Variance)*
Chronological age	0 218	5[a]	5[a]
Nonverbal intelligence	0.388	15[b]	13[b]
Nonword repetition	0 525	27[b]	15[b]
Sound Mimicry	0.295	9[b]	0
Total	0.578	33[b]	–

[a]$P < 0\ 05$; [b]$P < 0.01$
Source: Gathercole and Baddeley (1989)

development, by looking at the relationship between our nonword repetition test and vocabulary size.

Our results are shown in Table 4.4, which shows that on starting school there is a clear correlation ($r = 0.492$) between nonword repetition and vocabulary. The correlation remains when the effect of nonverbal intelligence is removed statistically. One year later the correlation remains high ($r = 0.572$), and remains statistically significant when the effect of vocabulary level on the previous test is removed, suggesting that the process underlying nonword repetition is continuing to be important for vocabulary learned during the first year at school.

These findings are clearly consistent with the view that the articulatory loop is central both to nonword repetition and to the acquisition of one's native language. However, it is important to remember that correlation does not necessarily mean causation. It is possible, for example, that both phonological STM and vocabulary learning are dependent on some third factor; phonological awareness, or the amount and richness of language that has already been learned, might be two possibilities. Indeed, such results make it clear that our understanding of the processes underlying the operation of the phonological loop is still at a very primitive stage. They do, however, suggest that they are of great potential importance.

Conclusion

We have discussed the phonological loop in some detail since it is the most extensively explored component of working memory. If this aspect of the enterprise fails, then it seems unlikely that it will succeed in the more complex problems of tackling the visuo-spatial sketchpad and the central executive.

How successful has the enterprise been then? It does seem to have provided a simple explanation of the abundant evidence suggesting that STM is in some sense a speech-based system. At the same time it clearly indicates that this is only one component of working memory. However, when looked at from a broader perspective, it appears that this component of working memory is potentially an important one for learning to speak and to read and for comprehending spoken discourse. As such, it is a system worth understanding.

How well do we understand it? At a qualitative level reasonably well, but in detail hardly at all. We know nothing about the nature of the store, its time characteristics, how information is read into it and retrieved from it, and how it relates to the processes involved in speech perception and production. We know relatively little about how it relates to phenomena and concepts based on techniques relying on reaction time measures, an area that is well reviewed by Monsell (1984). At the moment, I am not at all convinced that we even have an adequate language for developing a detailed model, although I have some hopes that the parallel distributed processing models described in Chapter 14 may offer one way of tackling the problem of modeling the phonological loop.

OVERVIEW

The chapter began by discussing a series of experiments concerned with the question of whether STS acts as a working memory, playing an important part in cognitive activities such as learning, comprehending and reasoning. The experiments used a dual-task approach in which working memory capacity was systematically absorbed by requiring subjects to hold sequences of digits in STS at the same time as they performed tasks involving reasoning, learning or comprehending. While clear impairment was found on these tasks, the extent and nature of the disruption was not as great as would be expected on the assumption that the same unitary system both holds digits and acts as a working memory. On the basis of these results, a multi-component working memory was proposed, with a controlling central executive system and a number of subsidiary slave systems.

One of these, the articulatory or phonological loop is described in more detail. It is assumed to comprise a short-term phonological store assisted by a control process based on articulatory rehearsal. It is shown that this simple model can account for a range of factors that influence memory span, including acoustic similarity, word length, unattended speech and articulatory suppression.

We finally considered the question of the role played by the phonological loop in everyday cognition. Evidence is presented to suggest that it plays an important role in learning to read, in the comprehension of language and in the acquisition of vocabulary; in all these areas, the evidence comes both from the development of language in normal children, and from the performance of patients suffering from impaired STM following brain damage.