Chapter 7

Mechanisms of Verbal Working Memory Revealed by Neuroimaging Studies

John Jonides

In 1970, when I had the privilege of beginning my work with Henry Gleitman, the study of cognition was in the midst of a vital period. The reasons for this vitality were many, but among them was the vision that perception, memory, language, and thinking could be understood by decomposing cognitive processes into their essential elements. The tools for this decomposition were chronometric analysis and the analysis of patterns of errors that subjects made in various tasks, both of which were being applied to a host of problems in the study of cognition. Henry and I embraced these developments and used these techniques to study processes of item-recognition and short-term memory in a series of studies that we conducted under the watchful eye of Henry and Lila Gleitman's Thursday evening (and often late-night) research seminar.

The strategy of understanding complex cognitive processes by decomposing them into their simpler parts continues to be widely taught and learned as an essential skill for students of cognition. At the time of my graduate education, implementing this strategy most often involved the careful design of behavioral experiments that would reveal underlying processes in the chronometric or errorful data of subjects. In the hands of clever psychologists such as Henry and Lila, this strategy yielded insights into a wide variety of cognitive processes. In the ensuing years, the strategy has been broadened while still maintaining its essence. The first significant broadening involved the widespread use of computer models as ways of further defining and refining conceptions of elementary processes and how these combine to yield cognition. The second significant broadening is the one that is the current focus of much interest among students of cognition: the use of neuroimaging techniques to reveal brain processes that mediate cognitive processes. These techniques allow us to go well beyond merely localizing processes to various brain areas; in combination with behavioral data on normal and brain-injured humans and in combination with suitable data from invasive studies of animals, neuroimaging studies can ac-

complish the same goal that we had in 1970: to decompose cognitive processes into their elementary components. In short, neuroimaging techniques provide another modality of data for understanding cognition.

To illustrate how the use of neuroimaging has enriched the study of cognition, I shall briefly review the results of a program of research conducted by Ed Smith, me, and our colleagues in recent years. We have been investigating what is now called working memory. The reason for focusing on working memory is simple: Various lines of evidence make it clear that working memory is an essential component of cognition in that it participates in such skills as problem solving, reasoning, categorization, and language comprehension. Indeed, an individual's working-memory capacity (if measured in the right way) has been found to correlate with a large variety of complex cognitive skills (see, e.g., Carpenter, Just, and Shell 1990; Daneman and Merikle 1996; Kyllonen and Christal 1990). Furthermore, it has been shown repeatedly that a decline in working memory with normal aging and with various brain pathologies is an important predictor of declines in performance in problem solving and reasoning, showing again that working memory is critical to cognitive life (see, e.g., Salthouse 1993). In short, the phenomena of working memory are examples of the sort of "big effects" that the Gleitmans advocated as worthy topics for intense research.

We define working memory in a canonical way. It is the memory system that keeps active a limited amount of information for as long as one continues to work with that information. The information in working memory is easily and readily accessible, and it is subject to frequent updating or substitution by new information. Finally, the feature that discriminates the concept of working memory from earlier conceptions of "short-term" or "primary" memory is that information held in working memory is subject to processing in various ways by what Baddeley (1986, 1992) and others have called executive processes. It is in this sense that working memory goes beyond mere storage; it involves the manipulation of information that is stored as well.

Our research group has devoted itself to several issues concerned with working memory; here we review research concerned with two of these. One concerns the architecture of the working memory system itself and whether this system consists of several subcomponents concerned with storage, rehearsal, and manipulation of information that is stored. On this issue, much of our research has concentrated on working memory for verbal information. A second issue is whether working memory is a single system or whether it consists of several subsystems each tied to the processing of different sorts of information.

The Architecture of Verbal Working Memory

In his influential model of working memory, Alan Baddeley (1986, 1992) proposed that working memory for verbal information consists of three components. One is a buffer responsible for the temporary storage of verbal codes. A second is a rehearsal mechanism that recirculates information in the verbal buffer for the purpose of preventing decay or interference. A third is a set of processing mechanisms, collectively called the central executive, that are capable of manipulating the information stored in the buffer (including rehearsing it). Evidence for a dissociation between storage and rehearsal processes has come largely from behavioral studies of normal and brain-injured subjects. Although this evidence is compelling, it is not decisive (see Jonides et al. 1996 for a review), and so we have conducted experiments using neuroimaging techniques to try to identify the subcomponent processes of working memory.

One experiment makes use of a paradigm that recruits all the components of working memory (storage, rehearsal, and executive), the "*n*-back" task. In this task, subjects are presented a series of single letters; as each is presented, subjects must decide whether it matches the one that was presented *n* items back in the series. In an experiment from our laboratory reported by Awh, Jonides, Smith, Schumacher, Koeppe, and Katz (1996), *n* was set to 2. A schematic of the memory condition of the task is shown in the top panel of figure 7.1. The panel illustrates that letters were presented for 0.5 sec. each with 2.5 sec. intervening between successive letters; subjects engaged in the memory task for a continuous period of approximately 60 seconds while they reclined in a PET scanner. Note that successful performance in this task requires storing in memory a constantly changing set of at least two letters, the "oldest" of which must be compared with the currently presented letter. Thus the task requires both storage of verbal information as well as processes that must update this information continuously.

Of course, performance in this memory condition also includes processes of perception and response, processes that were not the targets of our interest in this experiment. The typical strategy for eliminating the effects of such processes on brain activations in a neuroimaging experiment is to test subjects in a second condition that includes only these ancillary processes, by hypothesis, and then to subtract the brain activations from such a control condition from the condition of interest. We recognize that this "subtraction" strategy has earned some well-deserved criticisms, and we address this issue below. Nevertheless, we followed this strategy by also testing subjects in a memory control condition, shown in the second panel of figure 7.1. In this condition,
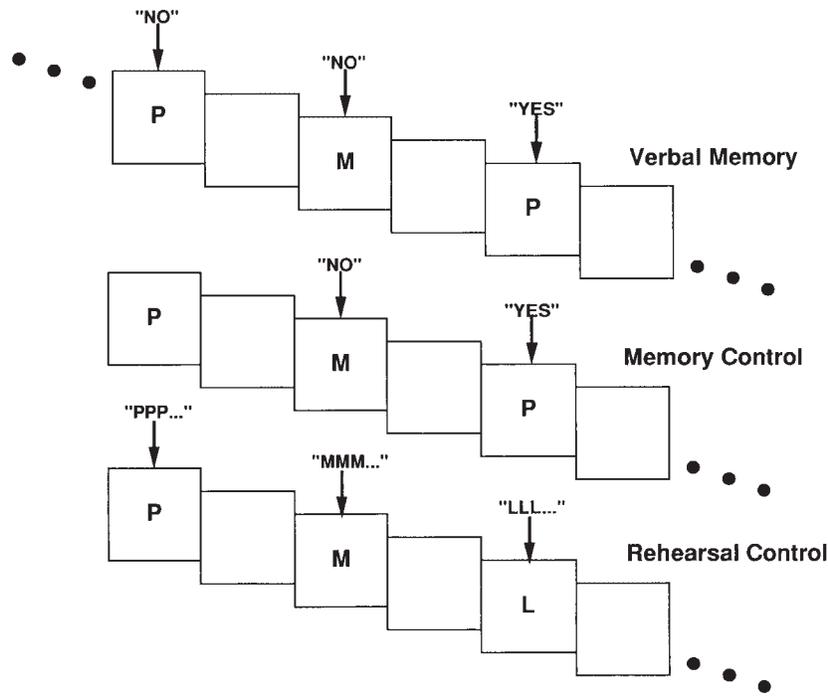
Figure 7.1.
Schematics of the three tasks used in the experiment by Awh et al. (1996).

subjects were presented with a series of letters, but they responded positively only if each letter matched a fixed target letter for which they searched on that trial (say, "P"). Activations from this control condition were subtracted from those in the memory condition. The results revealed activation in anterior parts of the cortex: in Broca's, supplementary motor, and premotor areas. In addition, there was activation in posterior parietal cortex in the superior parietal lobule and supramarginal gyrus.

The anterior sites that were activated in this task have been claimed to be part of a circuit (including the cerebellum, which also showed activation) that is responsible for rehearsal (Paulesu, Frith, and Frackowiak 1993); indeed, this is quite sensible in that some or all of these sites are involved in the production of explicit speech as well. The posterior sites have been implicated in storage and selective attention processes required in this task. How can one confirm these putative functions? Paulesu et al. (1993) provided some evidence for the rehearsal function of the anterior sites by showing that these sites were

also activated in another task in which judgments of rhyming were required, judgments that presumably also engage internal language-production mechanisms (see the chapter by Reisberg in this volume for additional evidence about such judgments). Furthermore, the site we found activated in the supramarginal gyrus of parietal cortex is the most common site of damage in patients who suffer deficits in verbal memory span (McCarthy and Warrington 1990), suggesting that it may be involved in storage of verbal material. Also, the site in the superior parietal lobule has been implicated as a region involved in shifts of attention from one item of material to another in many tasks, as shown by a parallel between this site and sites that are activated in tasks in which explicit shifts of attention are required (see Awh and Jonides 1999 for a review). So there is circumstantial evidence for the functions that we claim.

Beyond this, Awh et al. (1996) included a second rehearsal control condition in their experiment that allowed them to test for the rehearsal function of the anterior sites. It is schematized in the third panel of figure 7.1. In this condition, subjects were again shown a sequence of individual letters, and they were instructed to rehearse each one silently until the next one appeared; then they were to rehearse the next one, and so on. This condition requires internal speech, and so subtracting the activation due to this condition from that due to the memory condition should diminish activation in anterior sites but not in posterior ones if the anterior ones in the Memory condition are, indeed, responsible for rehearsal. In fact, this is just what Awh et al. (1996) reported. Thus it appears that experiments using neuroimaging technology can provide corroborating and converging evidence for claims rooted in behavioral data as well as provide information on the brain areas that are the substrates of working memory.

Note that the experimental rationale underlying the application of neuroimaging techniques owes a great debt to the experimental rationale first articulated by Donders (1868) for the study of reaction times, a rationale discussed and debated at great length in our research seminars some one hundred years after Donders described it. Donders argued that if one could construct two tasks such that the first contained all the processing components of the second plus an additional one, then the difference in reaction time between the two tasks should be a relatively pure measure of the time required for the additional process required by the first task. In a similar fashion, much neuroimaging research has relied on subtracting the activation of one or more control tasks from an experimental task of interest, as illustrated above, to reveal activation due to processes of interest that are required by the experimental task. Just as the rationale due to Donders has been called

into question, so also can one raise questions about the validity of assuming that the activations due to selective processes can be subtracted out of a neuroimaging experiment without affecting activations due to other processes (see Jonides et al. 1997). The issue here is much the same one that occupied us in the 1970s: finding a way to isolate in data the effects of certain variables that have an impact on experimental performance. In the 1970s the data that concerned us on this score were reaction times and errors; in neuroimaging research they are patterns of brain activation that result from some task.

One experimental strategy that goes beyond the subtraction method relies on parametric variation of some variable of interest (in some ways, mimicking the strategy first applied by Sternberg (1966, 1969) to the analysis of chronometric data). We have implemented this strategy for the study of the components of verbal working memory in several experiments (Jonides et al. 1997; Braver et al. 1997; Cohen et al. 1997). Our work is based on the paradigm illustrated in figure 7.2. The figure reveals various versions of the *n*-back task, in which *n* is varied systematically. In the most demanding version, shown at the top of the figure, subjects must decide whether each letter matches in identity the one that appeared 3-back in the series. We also included 2-back and 1-back versions in other conditions, and we included a 0-back condition similar to the condition used by Awh et al. (1996) as a control; in the 0-back condition, subjects were provided a single letter-target at the beginning of a trial, and they were to respond positively any time that letter appeared anywhere in the series that was presented on that trial. The design shown in figure 7.2, therefore, implements a systematic variation of task load in a working memory task. We collected PET scans for each of the conditions shown in figure 7.2, and we also collected PET data for a baseline control condition not shown in the figure in which a series of letters was presented, and subjects simply responded with a keypress when each letter was shown; there was no memory requirement at all. This baseline control condition served as a way to subtract from each of the memory conditions the activation that was due to idiosyncratic differences in brain activity among subjects. The main comparisons in the experiment, however, did not rely on subtraction methodology; rather they relied on a comparison of activation in various regions across variations in task load.

Figure 7.3 reveals that, as expected, increases in task load produce a decrement in performance. This decrement appears as an increase in response time as well as a decrease in accuracy. The decrement in performance is accompanied by strikingly parallel changes in brain activation with task load at each of several sites, as shown in figure 7.4. The data in this figure were accumulated by taking regions of activation and deacti-
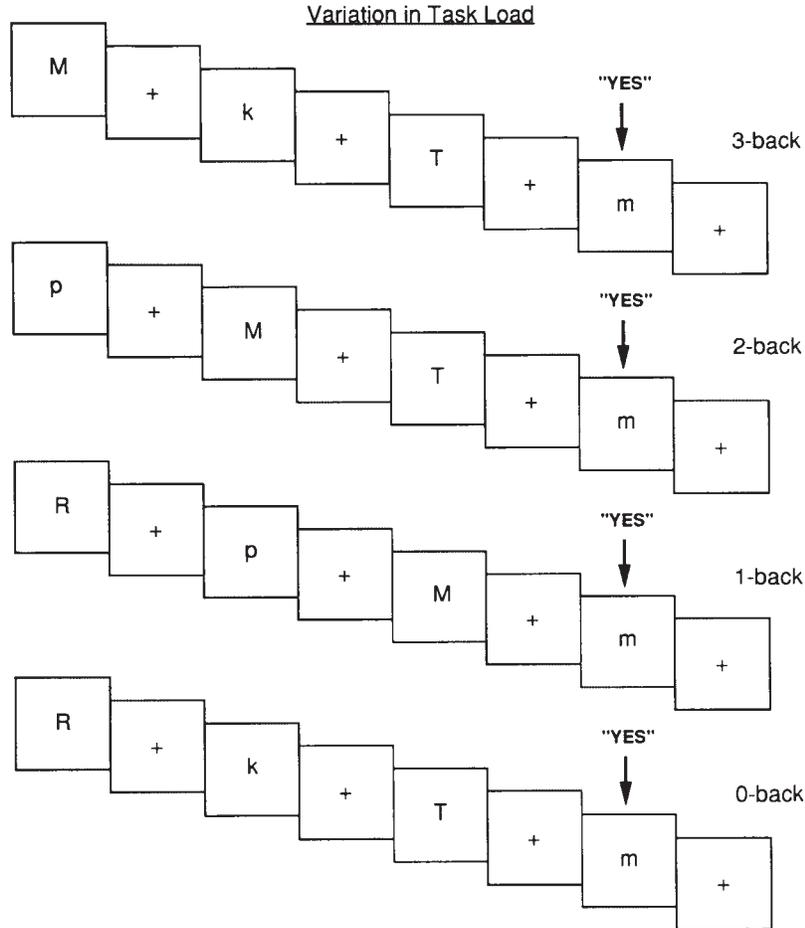
Figure 7.2.
Schematics of the tasks used by Jonides et al. (in press).

vation (that is, where the control condition shows less activation than the experimental condition and where it shows more activation, respectively) that had been previously identified with verbal working memory tasks (from the studies of Awh et al. 1996 and Schumacher et al. 1996). We then found levels of brain activation in the present data for each of these regions. The average activation in each of these regions was then plotted as a function of task load, and this is what is displayed in figure 7.4. The most striking feature of the data in this figure is that there is an overwhelming tendency for brain activation to increase
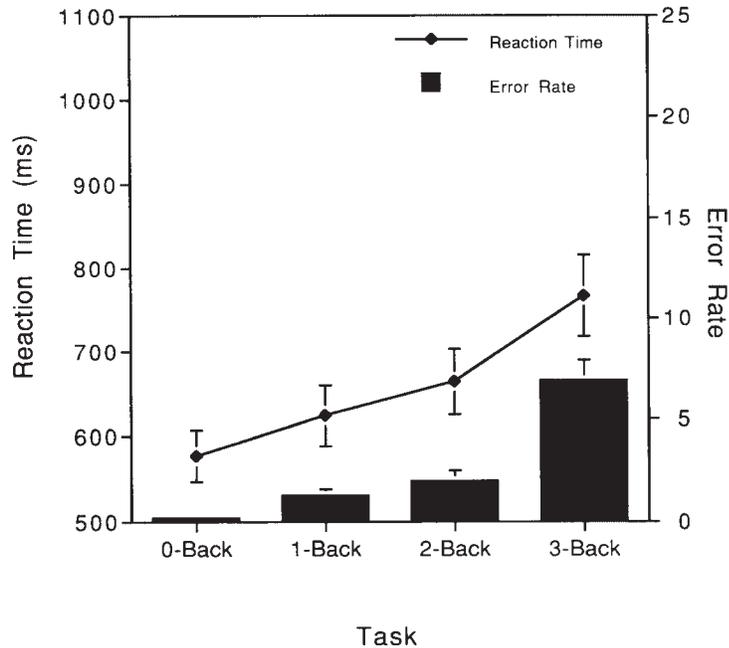
Figure 7.3.
Reaction time and error data from the experiment of Jonides et al. (in press) plotted as a function of memory load.

monotonically as task load increases, and for deactivation to decrease monotonically with task load (confirmed by statistical analysis: Jonides et al. 1997). Thus there is reason to suspect that the activations and de-activations that are shown in figure 7.4 are systematically related to the memory requirements of the task. Note that these changes occur in many regions, as the functions in figure 7.4 reveal. The presumption is that increases in activation reflect increased brain activity that is required by a task; likewise, decreases in brain activation in selected regions reflect requirements for inhibition of brain activity in those regions. The specific account of which regions increase and which decrease is beyond the scope of this review. Here we merely highlight that brain activation in many regions is systematically related to the requirements of the task.

Of course, one might argue that the variation in brain activation with task load shown in figure 7.4 is merely a reflection of increased overall effort as task load increases, and not selectively related to processes having to do with the memory requirements of the tasks per se. This argument is laid to rest by examining other areas of the brain that should
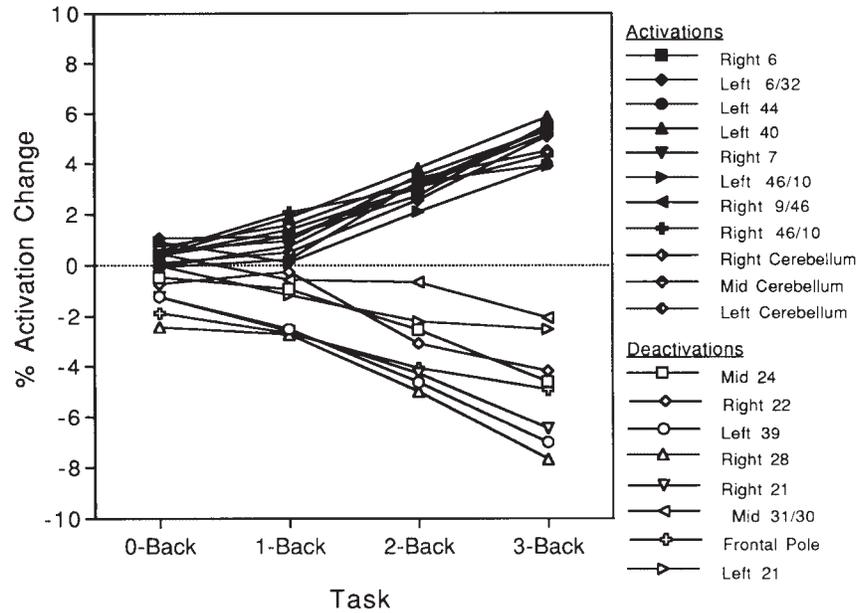
Figure 7.4.
Brain activations (above the horizontal line) and deactivations (below the line) for memory-related areas plotted as a function of memory load from the experiment of Jonides et al. (in press). Each function corresponds to one region of activation (labeled by its number according to the system of Brodmann or by anatomical structure). The identified regions are those that showed significant activation in a previous study using a similar task (see text).

not be recruited by memory processes: occipital areas that are involved in visual processing, somatosensory areas that are not relevant to the task, and primary motor areas whose activation should not vary with task load. The activations for these regions as a function of task load are displayed in figure 7.5, which shows that there is no systematic variation in brain activation in these regions as the memory task increases in difficulty. The contrast between the functions in figure 7.5 and those in figure 7.4 suggests that the areas identified in figure 7.4 do, indeed, reflect memory-sensitive processes that are active during these tasks.

Note that the outcomes of parametric studies of working memory validate findings reported using the subtraction methodology based on the logic of Donders. That is, the various regions of activation shown in figure 7.4 and in other studies using parametric techniques are just the ones, by and large, that show significant activation in subtraction paradigms such as the one described above from our laboratory (Awh et al.
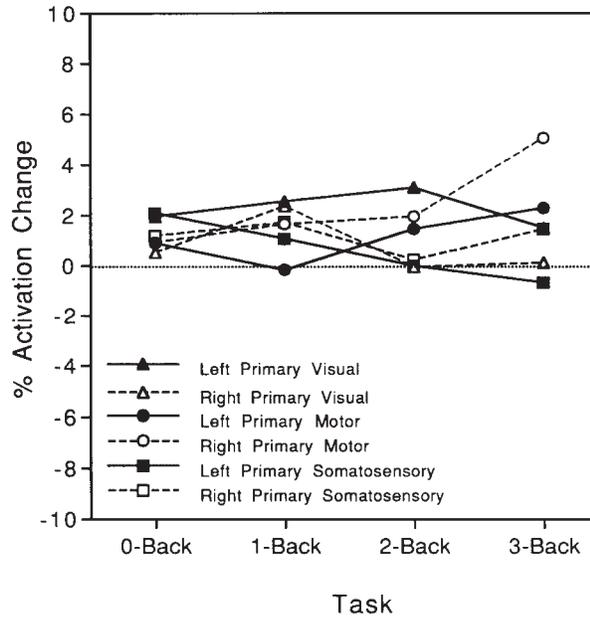
Figure 7.5.
Brain activations for motor, visual, and somatosensory areas plotted as a function of memory load from the experiment of Jonides et al. (in press). The areas shown in the figure were identified by placing regions-of-interest on primary visual, motor, and somatosensory areas of the brain and calculating the activations in those regions.

1996). Thus, although one must exercise caution in using subtraction logic and in interpreting the results of such experiments, the outcomes of experiments on working memory that have used this logic seem to be replicated in experiments with parametric experimental manipulation of relevant variables.

Another contribution of the parametric method is that it provides an opportunity to examine details of the "dose-response" curves that result, such as those shown in figure 7.4. We have exploited this property of parametric designs both in the experiment described above and in experiments using this paradigm with functional MRI as the imaging modality (Braver et al. 1997; Cohen et al. 1997). Functional MRI measurements permit a more detailed examination of brain activation patterns because they provide somewhat greater spatial resolution as well as an opportunity to examine the temporal dynamics of processes within a single experimental trial.

Consider, for example, an experiment in which we varied memory load in an *n*-back task while stretching out the retention interval on each
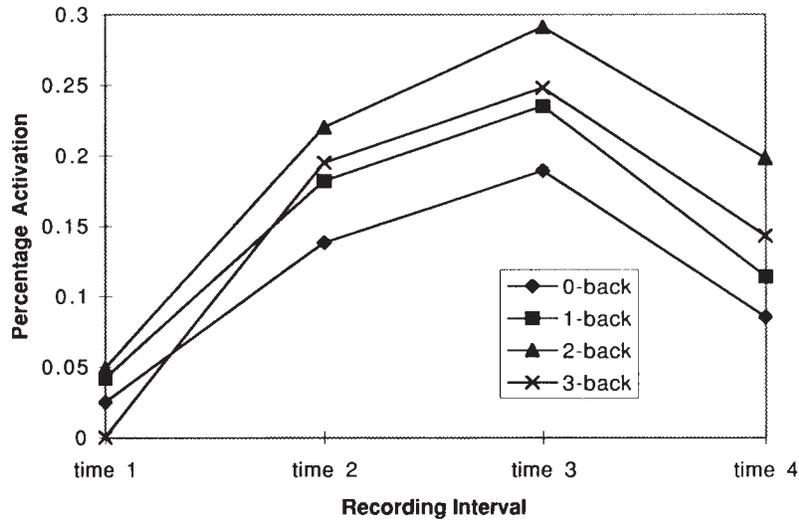
Figure 7.6.
Activations from a region of extrastriate occipital cortex as a function of the time of recording during retention interval (from the experiment of Cohen et al., in press).

trial so that we could examine the dynamics of activation in various brain regions during that retention interval (Cohen et al. 1997). The design is much like that displayed in figure 7.2, but the delay interval between successive letters in each condition was increased to 10 seconds. This permitted us to collect four scans of the entire brain, each one occupying a 2.5-second interval during the time between letters. Thus level of brain activation could be assessed four times during the interval when subjects were engaged in the working memory task.

   The value of this technique is revealed by examining the results of the experiment. Examine first figure 7.6, which shows the activations that were obtained in a region of extrastriate occipital cortex. The four functions in the figure correspond to the four conditions of memory load (0-, 1-, 2-, and 3-back). Each function has four points, each corresponding to an activation for one of the four recording intervals during the retention period. Note the four functions lie fairly close to one another and that there is little systematic effect of memory load. However, the four functions are also all noticeably bowed. This can be taken to mean that as time passed during the retention interval, the amount of activation increased in this region of occipital cortex and then declined. Why should this be so? The most reasonable interpretation is that these functions all reveal the activation that was caused by presentation of a stimulus letter. The activation rises over the course of the first 7.5 seconds of the
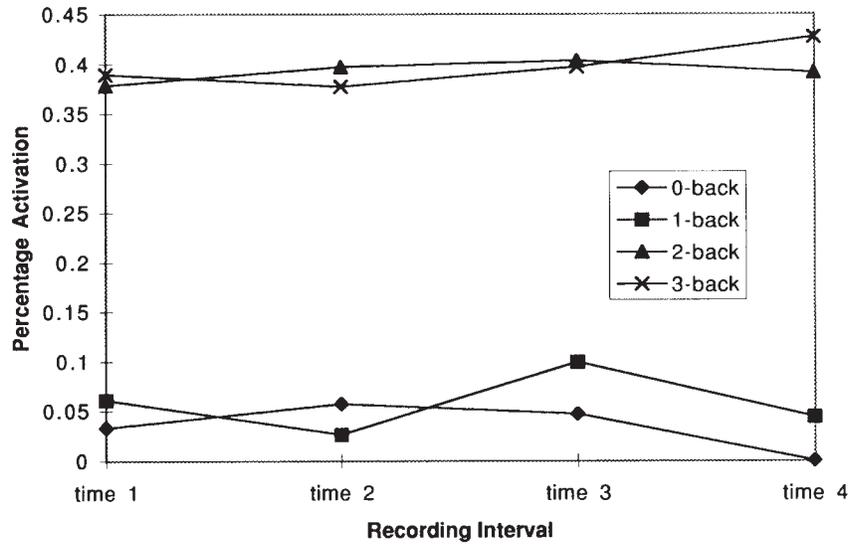
Figure 7.7.
Activations from a region of dorsolateral prefrontal cortex as a function of time of record-
ing during the retention interval (from the experiment of Cohen et al., in press).

retention interval because the hemodynamic signal that is recorded by
fMRI (corresponding to the neural signal that is tied to the increase in
blood flow) is delayed relative to the neural event that causes it. By this
interpretation, the functions in figure 7.6 reveal that fMRI recordings
can be sensitive to early processes of encoding, and that these processes
are not particularly sensitive to the memory load of the task.

Examine now the activations in figure 7.7. This figure shows func-
tions analogous to those of figure 7.6, but for activations recorded in a
region of dorsolateral prefrontal cortex in the right hemisphere. Note
that the pattern of data in this figure differs markedly from that in figure
7.6. First of all, the functions are not at all bowed in shape. That is, acti-
vations for all levels of memory load appear to be quite steady through-
out the retention interval, not transient, as they are in occipital cortex.
Notice also that in prefrontal cortex there is a dramatic effect of memory
load on activation level. As the figure indicates, the 2-back and 3-back
tasks resulted in substantially more activation throughout the retention
interval than did the 0-back and 1-back tasks. The sustained nature of
this activation and its sensitivity to memory load suggest that this area
of prefrontal cortex is somehow involved in the maintenance of repre-
sentations in working memory. This may be via direct storage of in-
formation or via some more indirect role. For example, regions in

prefrontal cortex may serve as pointers to other regions of the brain (possibly posterior regions) that are the sites of information storage. Whatever the specific role of these prefrontal structures, it is clear from data such as those of figure 7.7 that the activation of prefrontal cortex during the *n*-back task is not transient in nature. It might have been so if prefrontal cortex were involved strictly in executive functions such as updating the contents of working memory or temporally tagging letters as they are presented. In both of these cases, one would have expected the activation to be transient in nature, quite different from the data that were obtained.

These data indicate the potential that neuroimaging studies have for extending our knowledge of cognitive mechanisms. In connection with the sort of behavioral data that one can gather from working memory studies, neuroimaging data are proving helpful in specifying the various component mechanisms that contribute to complex cognitive phenomena.

Differing Subsystems of Working Memory

A central issue in the study of working memory has been whether it consists of a unitary processing system or whether it is composed of several subsystems. One line of evidence that has been illuminating about this issue comes from studies of brain-injured patients. There is evidence, reviewed in detail by Jonides et al. (1996), among others, that there are multiple subsystems of working memory defined by the type of information that is maintained. For example, there are reports of patients who have deficits in verbal working memory with no deficits in working memory for spatial information; by complement, there is a report of a patient who has a deficit in spatial but not verbal working memory.

In support of this distinction, there is also evidence from strictly behavioral studies of a dissociation between subsystems of working memory defined by the type of information that is processed. The behavioral technique that has been used to provide evidence for this claim involves experiments in which subjects engage in one or another working memory task while a second (presumably interfering) task is performed. The logic of these experiments is this: If secondary tasks can be found that require the use of one or another internal code for information that is processed, then they should selectively interfere with a primary working memory task to the extent that that primary task makes use of the same code. For example, a secondary task that engages a phonological code should interfere with verbal working memory if verbal working memory also requires the use of a phonological code; but it

should not interfere with spatial working memory if that subsystem uses a code that is not phonological. Likewise, a secondary task that makes use of a spatial code should interfere with spatial working memory but not verbal working memory. Various experiments have implemented this rationale (see, e.g., Meudell 1972; Salthouse 1974; Logie, Zucco, and Baddeley 1990; Logie 1995), and they have led to the view that working memory is, indeed, composed of several subsystems defined by the type of information that is processed.

These lines of neuropsychological and behavioral evidence are not immune to criticism, however. Experiments on neuropsychological populations about the dissociation of different working memory subsystems are limited to precious few patients, who are often tested on tasks that may not purely recruit one or another working memory subsystem. As for selective interference experiments testing normal subjects, it is often difficult to justify the assumption that a secondary task has a truly selective interfering effect on a primary task. There may be several sites of interference (see Jonides et al. 1996 for a detailed discussion). Consequently, we sought a line of evidence from neuroimaging studies of working memory that might converge with the behavioral evidence to address the issue of whether working memory is best conceptualized as a *set* of subsystems rather than a single system of information processing.

The details of our experiments are reported elsewhere (Jonides et al. 1993; Smith and Jonides 1995), but it is instructive to examine one example to see how neuroimaging evidence can strengthen the case for separable working memory subsystems (Smith, Jonides, and Koeppe 1996). Consider the pair of tasks illustrated in figure 7.8. Each panel shows schematics of the events on typical trials of the spatial and verbal memory conditions respectively; both conditions involve a 3-back task, with the nature of the memorandum differing between conditions. In the spatial memory condition, subjects saw a stream of letters with the letters appeared in varying locations on the screen. The subjects' task was to answer positively (via a button-response) if a letter's location matched the location of the one that appeared three back in the sequence; if not, they were to answer negatively (via another button-response). Similarly, the verbal memory condition also required subjects to match the current stimulus to the one that was 3-back in the sequence; however, in this case they were to match letters on their identities regardless of their spatial locations. Thus, in both tasks, subjects had to keep in working memory information about several previous stimuli, they had to match the current stimulus against the one that appeared 3-back in the sequence, and they had to update the contents of their memories with each succeeding stimulus presentation. The major
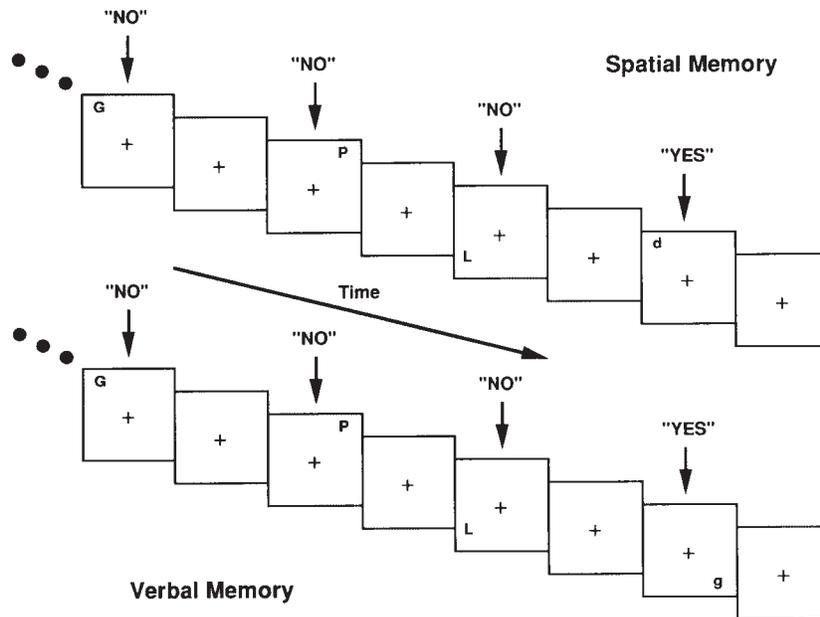
Figure 7.8.
Schematics of the memory tasks used by Smith, Jonides, and Koeppe (1996).

difference between the conditions was in the sort of information that was stored in memory: In the spatial memory condition it was locational information, and in the verbal memory condition it was identity information (storing a visual code for each letter would not suffice in the verbal memory condition because the case of the letters haphazardly varied from upper to lower).

We collected data for each of these two memory conditions while subjects reclined in a PET scanner. The scanning recorded all the brain areas that were activated during any portions of these tasks. To focus on the processes specifically involving memory, we also tested subjects in control conditions whose activations were then subtracted from those of the memory conditions. In the spatial control condition, subjects were shown three locations on the screen prior to a sequence of stimuli such as those shown in the top panel of figure 7.8, and they were to respond positively anytime any letter appeared in any of these positions; otherwise, they responded negatively. Likewise, in the verbal control condition, they were shown three letters, and they responded positively anytime any of these appeared.

The brain activations that resulted from the subtractions of the control from the memory conditions revealed both substantial overlap in

activations between conditions and substantial differences. Generally speaking, there was bilateral activation in both spatial and verbal tasks in both anterior and posterior regions of the brain. Beyond this, though, there was evidence that the spatial task activated some structures in the right hemisphere more than in the left; in a complementary way, the verbal task activated regions in the left hemisphere more than in the right.

Other than this noticeable difference between the tasks, there was considerable similarity in the regions activated. There was clear activation in two regions of posterior parietal cortex, one more lateral than the other, similar to some of the results described above. Also, there was activation in dorsolateral prefrontal cortex, also similar to results presented above. And there was evidence of activation in inferior frontal gyrus in the left hemisphere in the verbal task, an indication of the involvement of verbal rehearsal in this task.

For present purposes, our interest in this experiment is in its demonstration of different patterns of activation as a function of whether the material to be retained was verbal or spatial in nature. Although the overall circuitry revealed in these two conditions was similar, there was, as noted above, a dissociable pattern of activation with spatial material engaging the right hemisphere more than the left and verbal material engaging the left more than the right. Note that this pattern was obtained in an experiment in which the physical stimuli were nearly identical in the two memory conditions, rendering it unlikely that the different patterns of activation were a function of perceptual processing. We conclude that the different patterns are instead a reflection of different underlying circuitry for spatial than verbal working memory. This result, then, confirms and extends the data from neuropsychological and behavioral experiments and adds currency to the hypothesis that working memory is best characterized as a set of subsystems, each responsible for the processing of different sorts of information.

A Closing Reflection

Experiments such as those I have described are in the forefront of the news in cognitive psychology these days. They are there because neuroimaging techniques make it seem as if the often vague and ephemeral constructs of psychological theory can now be displayed in neural tissue. There is a certain excitement in being able to palpate something that was previously only imaginable, to see a functioning process on a computer display of brain activation where previously that process was only inferred from patterns of reaction times and errors. This is a lively development in the science. However, there are reasons beyond this to

value data from neuroimaging laboratories, and these reasons are mere extensions of the ones that guided the discussions in the Gleitmans' Thursday evening research seminars (which the clock suggested were endless, but which ended all too soon each week). Neuroimaging techniques can be applied, as our research suggests, to the identification of components of cognition and to the detailed description of the architecture of these components. In this way, it is valuable to conceive of neuroimaging data as an additional modality of insight into the phenomena of cognition, one that can supplement and enhance the behavioral study of normal and brain-injured humans.

Acknowledgment

References

Awh, E. and Jonides, J. (1999) Spatial selective attention and spatial working memory. In *The Attentive Brain,* ed. R. Parasuraman. Cambridge, MA: MIT Press.

Awh, E., Jonides, J., Smith, E.E., Schumacher, E.H., Koeppe, R.A., and Katz, S. (1996) Dissociation of storage and rehearsal in verbal working memory: Evidence from PET. *Psychological Science* 7:25–31.

Baddeley, A. D. (1986) *Working Memory.* Oxford: Oxford University Press.

Baddeley, A. D. (1992) Working memory. *Science* 255:556–559.

Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., and Noll, D. C. (1997) A parametric study of prefrontal cortex involvement in human working memory. *NeuroImage* 54:49–62.

Carpenter, P. A., Just, M. A., and Shell, P. (1990) What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices Test. *Psychological Review* 97:404–431.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., and Smith, E. E. (1997) Temporal dynamics of brain activation during a working memory task. *Nature* 386:604–608.

Daneman, M. and Merikle, P. M. (1996) Working memory and language comprehension: A Meta-Analysis. *Psychonomic Bulletin and Review* 3:422–433.

Donders, R. C. (1868) Over de snelheid van psychische processen. *Onderzoekingen gedaan in het psyiologish laboratorium der Utrechtsche Hoogeschool: Tweede Reeks* 2: 92–120.

Jonides, J., Reuter-Lorenz, P., Smith, E. E., Awh, E., Barnes, L., Drain, M., Glass, J., Lauber, E., Patalano, A., and Schumacher, E. H. (1996) Verbal and spatial working memory. In *The Psychology of Learning and Motivation,* ed. D. Medin.

Jonides, J., Schumacher, E. H., Smith, E. E., Lauber, E. J., Awh, E., Minoshima, S., and Koeppe, R. A. (1997) Verbal-working-memory load affects regional brain activation as measured by PET. *Journal of Cognitive Neuroscience* 9:462–475.

Jonides, J., Smith, E. E., Koeppe, R. A., Awh, E. Minoshima, S., and Mintun, M. A. (1993) Spatial working memory in humans as revealed by PET. *Nature* 363:623–625.

Kyllonen, P. C. and Christal, R. E. (1990) Reasoning ability is (little more than) working-memory capacity?! *Intelligence* 14:389–433.

Logie, R. H. (1995) *Visuo-spatial Working Memory.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Logie, R. H., Zucco, G. M., and Baddeley, A. D. (1990) Interference with visual short-term memory. *Acta Psychologica* 75:55–84.

McCarthy, R. A. and Warrington, E. K. (1990) *Cognitive Neuropsychology: A Clinical Introduction*. San Diego: Academic Press.

Meudell, P. R. (1972) Comparative effects of two types of distraction on the recall of visually presented verbal and nonverbal material. *Journal of Experimental Psychology* 94:244–247.

Paulesu, E., Frith, C. D., and Frackowiak, R. S. J. (1993) The neural correlates of the verbal component of working memory. *Nature* 362:342–344.

Salthouse, T. A. (1974) Using selective interference to investigate spatial memory representations. *Memory and Cognition* 2:749–857.

Salthouse, T. A. (1993) Influence of working memory on adult age differences in matrix reasoning. *British Journal of Psychology* 84:171–199.

Schumacher, E. H., Lauber, E., Awh, E., Jonides, J., Smith, E. E., and Koeppe, R. A. (1996) PET evidence for an amodal verbal working memory system. *NeuroImage* 3:79–88.

Smith, E. E. and Jonides, J. (1995) Working memory in humans: Neuropsychological evidence. In *The Cognitive Neurosciences*, ed. M. Gazzaniga. Cambridge, MA: MIT Press, pp. 1009–1020.

Smith, E. E., Jonides, J., and Koeppe, R. A. (1996) Dissociating verbal and spatial working memory using PET. *Cerebral Cortex* 6:11–20.

Sternberg, S. (1966) High-speed scanning in human memory. *Science* 153:652–654.

Sternberg, S. (1969) Memory-scanning: Mental processes revealed by reaction-time experiments. *American Scientist* 57:421–457.