COGNITIVE PSYCHOLOGY **33,** 5–42 (1997)
ARTICLE NO. CG970658

Working Memory: A View from Neuroimaging

Edward E. Smith and John Jonides

University of Michigan

We have used neuroimaging techniques, mainly positron emission tomography (PET), to study cognitively driven issues about working memory. Two kinds of experiments are described. In the first kind, we employ standard subtraction logic to uncover the basic components of working memory. These studies indicate that: (a) there are different working-memory systems for spatial, object, and verbal information (with the spatial system localized more in the right hemisphere, and the verbal system more in the left hemisphere); (b) within at least the spatial and verbal systems, separable components seem to be responsible for the passive storage of information and the active maintenance of information (with the storage component being localized more in the back of the brain, and the maintenance component in the front); and (c) there may be separate components responsible for processing the contents of working memory (localized in prefrontal cortex). In our second kind of experiment we have focused on verbal working memory and incrementally varied one task parameter—memory load—in an effort to obtain a more fine-grained analysis of the system's operations. The results indicate that all relevant components of the system show some increase in activity with increasing memory load (e.g., the frontal regions responsible for verbal rehearsal show incremental increases in activation with increasing memory load). In contrast, brain regions that are not part of the working-memory system show no effect of memory load. Furthermore, the time courses of activation may differ for regions that are sensitive to load versus those that are not. Taken together, our results provide support for certain cognitive models of working memory (e.g., Baddeley, 1992) and also suggest some distinctions that these models have not emphasized. And more fundamentally, the results provide a neural base for cognitive models of working memory.

Within the study of cognition there has always been a focus on a memory system that keeps active a limited amount of information for a brief period of time (e.g., Broadbent, 1958; Waugh & Norman, 1965; Atkinson & Shiffrin, 1968). Once referred to as ''short-term memory,'' this system is now more often called ''working memory.'' It is considered critical because it presumably serves as a mental blackboard for computations used in higher-level

---

The research reviewed in this article was supported by grants from the McDonnel–Pew Program in Cognitive Neuroscience, the National Institute of Aging, the Office of Naval Research, and the Department of Energy.

Address reprint requests to Edward E. Smith, Department of Psychology, University of Michigan, 525 E. University Ave., Ann Arbor, MI 48109-1109.

0010-0285/97 $25.00

processes such as reasoning, problem solving, and language understanding (e.g., Baddeley & Hitch, 1974; Baddeley, 1986; Just & Carpenter, 1992).

The vast majority of research on working memory has been carried out solely using behavioral experiments. This is clearly true of the many behavioral studies with normal subjects, and it applies as well to behavioral experiments with brain-injured patients that are primarily concerned with functional dissociation, with little focus on the specific brain regions involved. In the past few years, however, a line of research has emerged that involves studying working memory at both behavioral and neural levels. This research uses neuroimaging techniques, such as positron emission tomography (PET), to determine which brain areas are active when people perform various working-memory tasks, and to determine which cognitive functions are mediated by these areas. This approach to working memory places novel constraints on cognitive theories of working memory, and it offers new proposals about how cognitive processes are implemented in brain mechanisms. We have adopted this dual-level approach, and in this paper we review a number of our experiments.

The paper is divided into five sections. The first provides some background about PET. In the second section, we review neuroimaging experiments that address a fundamental question about the architecture of working memory: Are there different working-memory systems for different kinds of information? The third section reviews neuroimaging studies that ask another basic question about working memory: Does each system include separate components for storage and rehearsal? The experiments reviewed in the second and third sections all rest on a subtraction method, i.e., a contrast between a pair of conditions, one of which includes the process of interest and the other of which does not. In the fourth section of the paper, we consider some recent neuroimaging studies that have replaced subtraction with a method in which a single task parameter is varied quantitatively. These experiments ask whether quantitative variations in basic working-memory components can account for quantitative variations in brain activation and behavioral performance. The fifth and final section of the paper summarizes our major points and raises some remaining issues.

PET AND THE SUBTRACTION METHOD

The Logic of PET

The logic underlying the use of PET to study cognitive processes is this:

1. A selective change in some cognitive function—e.g., briefly remembering verbal information—is mediated by changes in neural activity in a region or in several selective regions of the brain;
2. A change in regional neural activity is accompanied by a change in blood flow to that region; and
3. Changes in regional blood flow may be monitored by monitoring the flow of a radioactive tracer in the blood stream using PET.

In brief, the collection of PET images involves several steps. Prior to the onset of a PET scan, a radioactive substance (typically 0–15 labeled water) is injected into a subject's bloodstream and flows to the brain regions that are neurally active during the task. As the substance decays in the brain, it emits positrons (a subatomic particle that is opposite in charge to an electron). Each positron moves only a few millimeters before it collides with a free-floating electron, leading to an annihilation of both particles. This annihilation process produces two photons that travel outward from the point of collision in approximately opposite directions. The photons are detected by crystals around the rings of a PET scanner that encompasses a subject's head. When two photons are detected on opposite sides of the detector within a very brief observation window, they are assumed to have come from the same annihilation process. Using tomographic techniques, it is then possible to construct images of where in the brain the annihilations occurred, and, by inference, which brain regions were active during the task.

The Logic of Subtraction

Returning to the general logic of PET, note that a comparison of experimental conditions is needed in order to observe a *change* in cognitive function. Following the pioneering work of Posner and his colleagues (e.g., Posner, Peterson, Fox, & Raichle, 1988), such comparisons have typically been based on the subtraction method. This method assumes that one can isolate the neural mechanisms corresponding to a specific cognitive process via the following logic, borrowed from behavioral studies using reaction-time measures (see Posner, 1978, for a review):

1. Construct a pair of behavioral tasks that differ in that one recruits the process of interest (the ‘‘target’’ task), whereas the other does not;
2. Obtain PET measures of regional cerebral blood flow during both tasks; and
3. Determine the difference in regional blood flow between the two tasks, inferring that regions unique to the target task mediate the cognitive process of interest.

To illustrate the logic of PET and the subtraction method, consider a hypothetical PET study of working memory. A subject participates in two conditions, which are blocked, and is injected with a radioactive tracer prior to performing each task. In a *memory* condition, on each trial a set of target letters is presented followed by a probe letter; the subject must decide whether or not the probe matches any of the targets, and indicates a decision by pushing one of two buttons. In the other condition, or control, the same materials are presented on each trial but there is no memory requirement, and the subject pushes either of the two buttons when a probe-letter appears. The control condition requires less working memory than the memory condition, but it is presumably comparable to the memory condition in all other respects.

In the memory condition, the neural networks that implement perceptual, response, and mnemonic processes are presumably activated; these activations lead to increases in blood flow to the areas comprising the networks, which is revealed by increased radioactivity in these areas. In the control condition, the neural network that mediates mnemonic process should be relatively quiescent, but the networks subserving non-memory processes should be as active as in the memory condition. Hence the subtraction of the control condition's activation pattern from that of the memory condition should reveal only the memory network of interest. This logic underlies all the experiments reported in the next two sections.

THE ARCHITECTURE OF WORKING MEMORY: MULTIPLE MEMORY STORES

A Cognitive Architecture of Working Memory

One of the most influential proposals about the cognitive architecture of working memory is due to Baddeley (1986; 1992). He assumes that working memory is composed of three major components: a buffer specialized for verbal material (the ‘‘phonological loop’’), another buffer specialized for visual–spatial material (the ‘‘visuospatial sketchpad’’), and a ‘‘central executive’’ that regulates the activities of the two buffers and utilizes the information contained in them. The phonological loop can be further divided into a pure storage buffer and a rehearsal process, with rehearsal responsible for recirculating decay-prone information in the buffer via internal speech. Baddeley and others have amassed a good deal of evidence for these architectural distinctions (see Baddeley, 1986, 1992, for reviews). Most of this evidence comes from behavioral experiments with normal subjects, but some derives from experiments with neurological patients. In all of these experiments, the dependent measures are exclusively behavioral. In the neuroimaging research that we report, the measures include biological ones (particularly, regional blood flow) as well as behavioral ones.

Much of our research about the architecture of working memory has been concerned with the distinctions featured in Baddeley's model. One key question is: Does the neural circuitry for working memory depend on the kind of stored information? We are concerned here not only with the possibility of different circuits for verbal and visual information, which is in keeping with the cognitive distinctions of Baddeley (1986, 1992), but also with possible differences in neural circuitry for visual–spatial and visual–object information, which goes beyond the distinctions articulated by Baddeley.

Phrasing the question of architecture in terms of neural circuitry does *not* imply that we are concerned primarily with the neural level. To the extent that we can document a neural distinction between verbal- and visual–spatial working memory, we will have provided additional justification for this assumption, justification that goes beyond that provided by behavioral evidence. Moreover, to the extent that we can document a neural distinction between

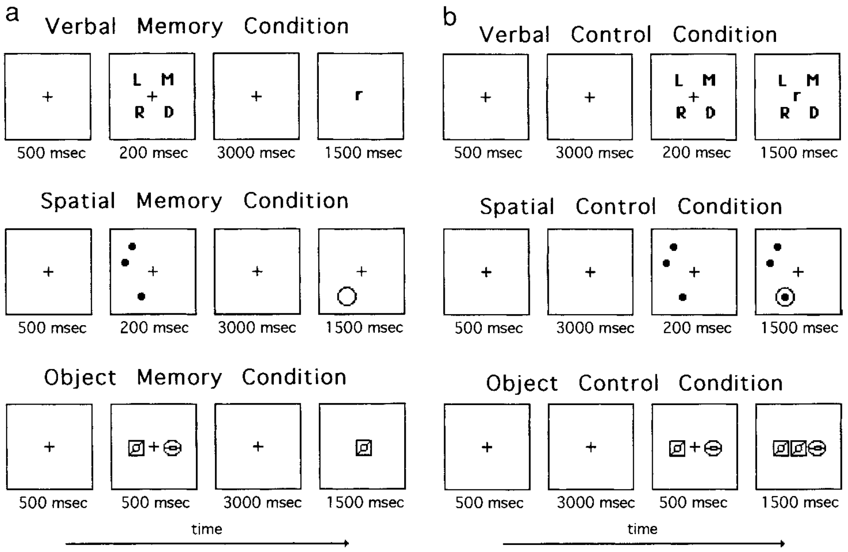

FIG. 1. Schematic representations of trials in three different working-memory tasks (a), and their corresponding controls (b). The top panel contains trials for the verbal tasks, the middle panel for the spatial tasks, and the bottom panel for the object tasks. All tasks are one-dimensional. (Adapted from Awh *et al.*, 1996; Jonides *et al.*, 1993; and Smith & Jonides, 1994.)

visual–spatial and visual–object working memory, we will have provided a reason for adding this distinction into the developing cognitive architecture.

Initial Experiments: One-Dimensional Studies

Basic tasks. Our initial studies used variants of Sternbergs' (1966) item-recognition paradigm and manipulated whether the information to be remembered was verbal (letters), visual–spatial (dots in various locations), or visual–object (novel geometric forms). The three kinds of memory tasks are illustrated in the three panels of Fig. 1a. In all tasks, the sequence of events on each trial was as follows. First, a fixation point was presented. Then followed the target information: four letters presented for 200 ms, or three dots presented for 200 ms, or two objects presented for 500 ms. (The variations in the number and duration of the items was intended to produce comparable behavioral accuracy across the three tasks.) Next came a 3000-ms. retention interval. Finally, a probe-recognition item appeared: a letter, a circle, or an object, respectively; for the letter or object probe, subjects had to decide whether or not it matched any of the target items; for the circle probe subjects had to decide whether it encircled any of the target locations. Subjects indicated their decisions by making one of two manual responses. Because only one source or dimension of information is available in any one task, we refer to these studies as ‘‘one-dimensional.’’

In keeping with the subtraction method, each of these memory conditions was paired with a control condition. The control conditions are presented in the three panels of Fig. 1b. In each case, the same kind and quantity of information is presented as in the memory condition, but the target information is still available when the probe is presented. Hence subjects can base their decisions strictly on a perceptual comparison, with no memory required.

Global and regional analyses. The tasks depicted in the three panels of Fig. 1 were conducted as separate experiments (the spatial-task results were initially reported in Jonides, Smith, Koeppe, Awh, Minoshima, & Mintun, 1993; the object-task results were reported in Smith & Jonides, 1994; and the verbal-task results were reported in Awh, Jonides, Smith, Schumacher, Koeppe, & Katz, 1996). It is most informative, though, to contrast the different results. Such contrasts are presented schematically in the three panels of Fig. 2. Each panel contains a representation of the left and right hemisphere, and each representation shows most of the significant cortical activations for two conditions—spatial versus verbal in the top panel, spatial versus object in the middle panel, and verbal versus object in the bottom panel. In each case, the activations depicted were determined by subtracting the activations from a control condition from the activations in the corresponding memory condition. The subtractions were first done on a pixel-by-pixel basis for each subject and then averaged across subjects.[1]

Consider first the difference in activations for the spatial versus verbal conditions (Fig. 2, top panel). The results are striking: all four areas activated for spatial memory are in the right hemisphere, whereas six of the seven cortical areas activated for verbal memory are in the left hemisphere and the seventh is in a midline structure (only five of these areas are presented in the schematic). At this *global level of analysis,* there is a qualitative double dissociation between spatial and verbal working memory, with the mechanisms needed to execute the two different tasks generally housed in different hemispheres.

One can learn far more about the processes of interest by moving to a *regional level of analysis.* This is an analysis that considers the specific neural areas involved and focuses on their known functions (as determined primarily by studies of brain-damaged patients, but also by prior neuroimaging experiments, and by single-cell recording studies with nonhuman primates). For the spatial task, two of the right-hemisphere regions activated are in the back of

---

[1] The analysis of PET images involves multiple steps. A typical scenario is as follows. First, an algorithim is used to correct for any subject movement between scans. Then the PET images for each subject are transformed to a standard stereotactic system (Talairach & Tournoux, 1988). Next a subtraction image is created for each subject between the average image for the target (memory) task and the average image for the control task. The subtraction images are next averaged across subjects, producing a group average subtraction-image; this group average subtraction-image consists of means and standard deviations of cerebral blood flow for each voxel. Finally, *t* statistic values are calculated for each voxel, correcting for mulitple comparisons (e.g., Friston, Frith, Liddle, & Frackowiak, 1991).

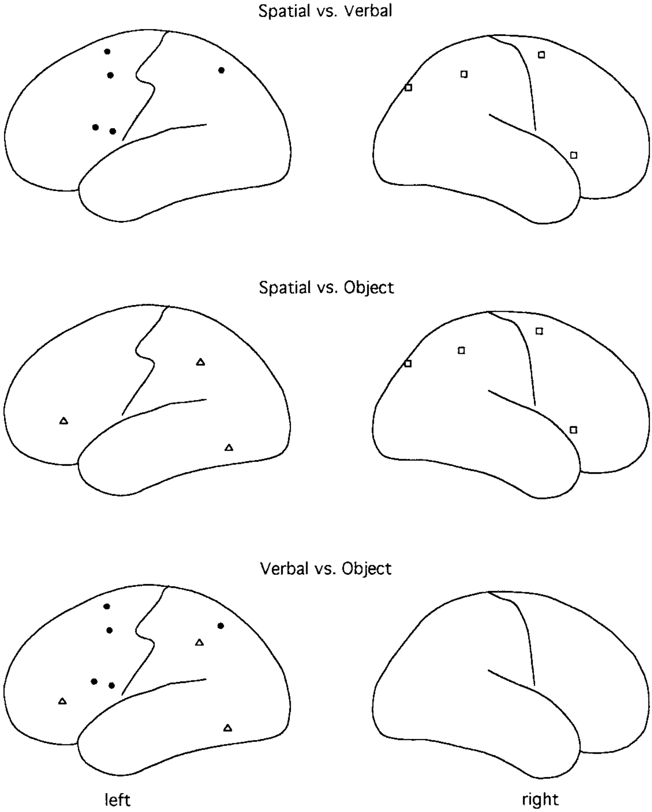

FIG. 2. Schematic representation of PET activations in left- and right-hemispheres for the one-dimensional tasks, with control activations subtracted. In the top panel the filled circles designate activations in the verbal task and the squares indicate activations in the spatial task. In the middle panel, squares again designate spatial activations, whereas triangles indicate object activations. In the bottom panel, filled circles indicate verbal activations and triangles objects activations. Some activated areas are not shown in the schematics because they were in a midline structure (the anterior angulate, which was activated in both the verbal and object tasks), or in subcortical regions (the left-hemisphere thalamus and right-hemisphere cerebellum, activated in the verbal task), or beneath the lateral surface of the cortex (left-hemisphere insular cortex, activated in the verbal task).

the brain, one is in the posterior parietal cortex, and the other is in the anterior occipital cortex (see Fig. 2). Studies of brain-damaged patients indicate that the parietal region is involved in spatial processing and spatial memory (e.g., McCarthy & Warrington, 1990). For example, patients with damage to their right-hemisphere posterior parietal region may be unable to decide whether the two segments of a bisected line are equal in length, or to maintain correct spatial relations between objects when copying a drawing (e.g., Kolb & Whishaw, 1995). As for the occipital region, other neuroimaging experiments suggest that this area is involved in the maintenance of visual images (e.g., Kosslyn, Alpert, Thompson, Maljkovic, Weise, Chabris, Hamilton, Rauch, &

Buananno, 1993; Kosslyn, Thompson, Kim, & Alpert, 1995). It is plausible, then, that in our spatial task, subjects computed the spatial positions of the target locations and used a parietal–occipital circuit to maintain an image of them during the retention interval. The functions of the two regions activated in the front of the brain in the spatial task are less clear. They may be involved in rehearsing the visual–spatial material (analogous to using implicit speech to rehearse verbal material—see below).

A regional analysis of the results of the verbal task suggests a two-component architecture, with structures in the back of the brain mediating storage, and structures in the front mediating rehearsal. Specifically, two of the most activated areas are in the left-hemisphere posterior parietal cortex (only one is shown in Fig. 2 because the activations were continuous). This is the region most likely to be damaged in patients who show a deficit in verbal short-term memory tasks. That is, there are neurological patients whose only detectable lesion is in the left-hemisphere posterior parietal area and whose major cognitive deficit is that their memory span for verbal materials is only one or two items (see Shallice, 1988). This strongly suggests that this region has a role in the storage or retrieval of verbal information. (Note that this region is the homologue of the area that we proposed as a spatial buffer, suggesting some symmetry between the neural mechanisms for spatial and verbal working memory.) Furthermore, all of the frontal regions activated in the verbal task—including Broca’s area, the premotor area, and the supplementary motor area (see Fig. 2)—are known to be involved in mediating higher-level aspects of speech (see Fuster, 1995). It is therefore likely that these frontal regions mediate an implicit-speech or verbal-rehearsal process in the verbal task.

The contrast between the spatial and object memory conditions is presented in the middle panel of Fig. 2. Again, a global analysis reveals a qualitative double dissociation, as three of the four areas of activation in the object task are in the left hemisphere while the fourth region is in a midline structure. A regional analysis of the activations in the object task again reveals the posterior parietal region and premotor region that were noted in the verbal task, plus two additional regions. One of these is an area in inferotemporal cortex, known to be involved in object recognition (e.g., Mishkin, Ungerleider, & Macko, 1983). This inferotemporal region may also play some role in the short-term storage of object information, but there are no obvious object-rehearsal regions activated. Indeed, the activations in the left-hemisphere premotor area (known to be involved in speech) and in the left-hemisphere posterior parietal region (the would-be verbal buffer) suggest that subjects sometimes verbally described the objects to themselves and then silently rehearsed these descriptions.

The last contrast, that between the verbal and object task (Fig. 2, bottom panel), reveals no new information at the regional level. At the global level, there is a double dissociation, but it is less pronounced than the other two. The inferotemporal activation is unique to the object condition, and there are a few frontal activations unique to the verbal condition. These results may underestimate the difference between verbal and object working memory because of the verbal coding that may have occurred in the object condition.[2]

In sum, these experiments provide evidence that spatial, verbal, and object working memory are mediated by different circuits. The findings support cognitive distinctions among three (at least) kinds of working memory. In addition, the known functionality of many of the specific areas activated are consistent with the broad functions that are required by the tasks.

Relevant Behavioral Findings

It is useful to consider some behavioral findings of the preceding studies to see how they relate to the activation results. There are four behavioral results that are of particular interest in any subtraction-based neuroimaging experiment. These amount to demonstrations that:

1. Accuracy is relatively high in all conditions, but not at ceiling levels; without such findings, one cannot be sure that subjects are executing the processes whose neural bases we are trying to measure.
2. Performance (e.g., accuracy and speed) should be better under the control conditions than in the corresponding target (e.g., memory) conditions; without such findings, there is no support for the subtraction logic.
3. Performance manifests some known ‘‘behavioral signatures’’ of the processes of interest, i.e., robust findings that have been obtained in prior behavioral studies; such findings help ensure that the processes measured during neuroimaging are the same as those studied during purely behavioral experiments.
4. There are correlations between behavioral performance and neural activation (as inferred from blood flow in PET); such findings help establish a direct link between the neural and cognitive levels.

Findings (1) and (2) are more critical than (3) and (4). No one would put much stock in a PET study of spatial working memory in which performance in the memory condition was at chance, and no one would place much faith in the logic of subtraction if performance under a control condition was significantly worse than that in a memory condition. Failing to find evidence of a correlation between behavioral evidence and PET data, by contrast, could

---

[2] To keep things manageable, we have not discussed every area that was significantly activated. Here we briefly note some of the areas that were not discussed. The object task activated the anterior angulate, which is likely involved in selective attention (e.g., Posner & Peterson, 1990). The verbal task also activated the anterior cingulate as well as the insular cortex, and two subcortical areas: left-hemisphere thalamus and right-hemisphere cerebellum. While the cognitive functions of these latter areas are far from clear, we note that the right cerebellum has a single-synapse connection to the left-hemisphere speech regions, and that it consistently shows up in studies of verbal working memory (see below). It may be part of the rehearsal circuit.

merely be a function of reduced range in behavioral performance engendered by the requirement that subjects perform well on the tasks.

How do the behavioral findings from the preceding studies fare against the four types of desired results? In all three experiments, accuracy under the memory conditions was relatively high but not at ceiling levels: 91% for the verbal task, 84% for the spatial task, and 80% for the object task (chance is 50% in all cases). This is hardly surprising, given that we used pilot testing to refine the tasks until they produced accuracies in the desired range.

Performance was also clearly better under the control than the corresponding memory condition in the spatial and object tasks, though not necessarily in the verbal task. For accuracy levels, the control vs memory contrast was 100 vs 84% for the spatial conditions, 94 vs 80% for the object conditions, and 91 vs 90% for the verbal conditions. We measured reaction times only in the verbal tasks: latencies averaged 724 ms in the memory condition vs 765 ms in the control condition.

Our spatial memory task was the only one that afforded a behavioral signature. The signature involved performance on distractor probes (i.e., probe circles that did *not* encircle a target location). Half of the distractors were relatively *near* one target location, whereas the other half were relatively *far* from all target locations. If subjects base their responses to the probes on a representation that is truly spatial, then performance should be better on the far than the near distractors. This near–far difference in accuracy has been obtained in strictly behavioral studies (Smith, Jonides, Koeppe, Awh, Schumacher, & Minoshima, 1995), and it was obtained in the PET experiment as well: 92% correct on far distractors versus 71% on near distractors. There were no comparable variations in probe similarity under the verbal and object memory conditions.

Finally, there is the matter of correlations between behavioral and activation measures. Only in the spatial-memory condition did such a correlation approach significance: there is a positive correlation of .33 between behavioral accuracy and the percentage increase in blood flow. We attribute the paucity of correlations to the lack of variability in behavioral performance, where this lack is the direct consequence of our insistence on relatively high accuracy.

In sum, the behavioral results of our experiments manifest the first three types of desired results. The only possible exception to this claim involved a violation of the subtraction logic in the verbal tasks (where the control was roughly as accurate as the memory condition and took longer); this potential exception will be addressed in the next experiments we review. Correlations between performance and activation remain elusive. Because roughly the same behavioral picture holds in the subsequent experiments that we report in this section and the next one, we will spare the reader many of the details.

Two-Dimensional Studies

Our one-dimensional studies are relatively pure with respect to information content; e.g., our spatial task provided information about only position. But

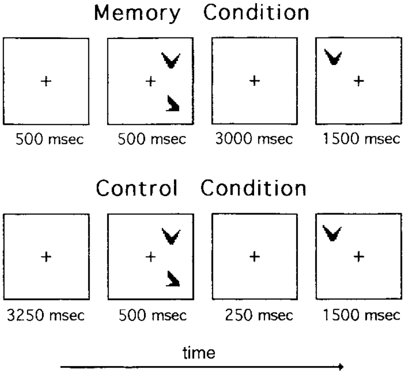

FIG. 3. Schematic representation of trials in the spatial and object two-dimensional tasks. The memory task is presented on top, and the control task on the bottom. (Adapted from Smith *et al.,* 1995.) For the memory trial shown, the spatial task would require a negative response and the object task a positive response.

these tasks have their drawbacks, one of which stems from their purity: the three tasks used different stimuli—letters vs dots vs objects—and it is possible that the different patterns of activation that we observed partly reflect specific stimulus effects. Another limitation of the one-dimensional tasks concerns their control conditions. The final event in a control trial consisted of the probe plus the targets, a display that differs substantially from the corresponding event in the memory conditions (probe only). This is a violation of the subtraction logic; it is the likely cause of the one potential behavioral violation of the subtraction method we observed (in the verbal task), and it may have contaminated some of the activation results.

Spatial versus object working memory. Our next pair of experiments circumvented the above problems by using *two-dimensional tasks.* Consider first the new task for contrasting spatial and object working memory, presented schematically in Fig. 3. The sequence of trial events was identical for the spatial and object tasks—only the instructions differed. Now the targets were two irregular objects (polygons) presented at random locations, and the probe consisted of a single object. The decision that subjects made about the probe depended on whether they were performing a spatial- or an object-memory task. In the spatial task, subjects decided whether the probe was in the same position as either target object; in the object task, subjects decided whether the probe was identical in form to either target object. The perceptual input was therefore identical in spatial and object tasks. As for control conditions, they were identical to the memory tasks except that the retention interval was reduced from 3000 to 250 ms (see Fig. 3). Now the processes required by each control condition should have been a proper subset of those needed in the corresponding memory condition (in accordance with subtraction logic) (See Smith *et al.,* 1995, for details).

The results of this experiment converged nicely with our initial findings.

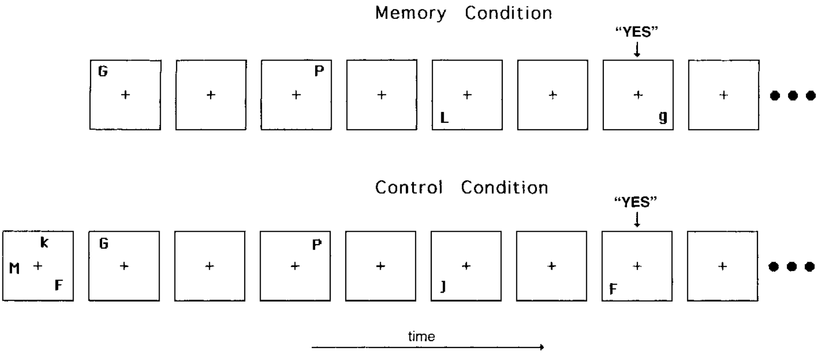

FIG. 4. Schematic representation of trials in the verbal and spatial two-dimensional tasks. The memory task is presented on top, and the control task is presented on the bottom. If these were trials in the verbal task, the correct response to the fourth item presented would be ''yes,'' as indicated. Note that letters vary in case, to force subjects to respond on the basis of letter identity rather than visual form. (Adapted from Smith *et al.,* 1996.)

At a global level of analysis, there was a double dissociation between spatial and object memory: after subtraction of the control condition, there were six areas activated in the spatial-memory task, two in the object-memory task, and only one of the latter appears among the former.

At a regional level, the object task activated two of the four regions that were activated in the initial task (left-hemisphere posterior parietal and inferotemporal areas). The spatial task activated all four regions that were activated in the initial task. It also activated a couple of regions which have been found in other neuroimaging studies of spatial working memory. One of the additional areas deserves special note. It is the left-hemisphere posterior parietal region, which is the supposed locus of the verbal/object buffer and the homologue of the right-hemisphere region that we earlier proposed to mediate spatial storage. This result suggests that, in contrast to our earlier findings, spatial working memory is not completely a right-hemisphere function (though there is still more activation in the right- than left-hemisphere posterior parietal region). This increase in bilateralization may reflect the greater difficulty of the two-dimensional task, or the fact that (for whatever reason) the task is intrinsically bilateral.

Spatial versus verbal working memory. The second experiment in this set contrasted spatial and verbal working memory. The task, presented in Fig. 4, is a new one. Instead of a series of discrete trials, subjects viewed a continuous stream of single letters, presented at randomly chosen locations, with a 2500-ms. interval between successive letters. Thus the location and identity of the letters varied, and which dimension the subject attended to depended on the instructions. In the letter-memory condition, subjects decided whether the identity of each letter matched that of the letter presented three back in the sequence (regardless of location) (see Fig. 4). Under the spatial-memory condition, subjects decided whether each letter matched the position of the letter presented three back (regardless of identity). Thus, similar to the preceding experiment, the sequence of perceptual events was identical under the two memory conditions, and all that differed was whether the subjects were instructed to store verbal or spatial information. Separate control conditions were used for the verbal and spatial conditions. Both controls are search tasks, and both involved the same sequence of events as under the memory conditions. In the verbal (spatial) control, subjects decided whether each item matched one or a fixed set of three letters (positions) that were identified at the beginning of the experiment. (See Smith, Jonides, & Koeppe, 1996, for details.)

Note that this ‘‘3-back task’’ differs from the item-recognition paradigm in that it requires coding the stored letters with respect to their temporal position (only a match 3-back counts, not one 1- or 2-back), and constantly changing these temporal codes as new letters are presented. The 3-back task therefore requires subjects to perform computations on information stored in working memory, which many researchers have taken as the paradigm case of working memory (e.g., Baddeley & Hitch, 1974; Daneman & Carpenter, 1980). Such computations would be expected to lead to increased activation in a particular region of frontal cortex, the dorsolateral prefrontal cortex, since this region is routinely activated in neuroimaging studies that require processing the contents of working memory (e.g., Cohen, Forman, Braver, Casey, Servan-Schreber, & Noll, 1994; Petrides, Alivisatos, Evans, & Meyer, 1993a; Petrides, Alivisatos, Meyer, & Evans, 1993b).

The results of this experiment are presented in Fig. 5. Rather than using schematic brains, it is worthwhile to present the current results in more detail. The PET activations are superimposed on a surface rendering of a single brain created from a standard MRI image. The results provide some convergence with our earlier contrast of verbal and spatial memory, but the activations are substantially more bilateral in the present case. At a global level, in the verbal memory task, the activation is clearly concentrated more in the left than in the right hemisphere (see Fig. 5). In the spatial task there is substantial activation in both hemispheres, though in key regions there is more activation in the right than in the left hemisphere (see below). This double dissociation, though quantitative rather than qualitative, is in line with our previous double dissociation between spatial and verbal working memory. It provides further evidence that two different systems are involved.

At a regional level of analysis, the verbal task activated two familiar, distinctly verbal regions: left-hemisphere posterior parietal cortex and a speech region, Broca’s area. There are also a few regions that are active in the current verbal task but were not in the initial task. These additional regions are no surprises, given the greater difficulty of the present task and the fact that it involves computations on working memory. One additional region is the right-hemisphere homologue of the left-hemisphere posterior parietal region. Just as we saw in the preceding study, then, some bilaterality appears once one moves away from a simple, one-dimensional task. The other addi-

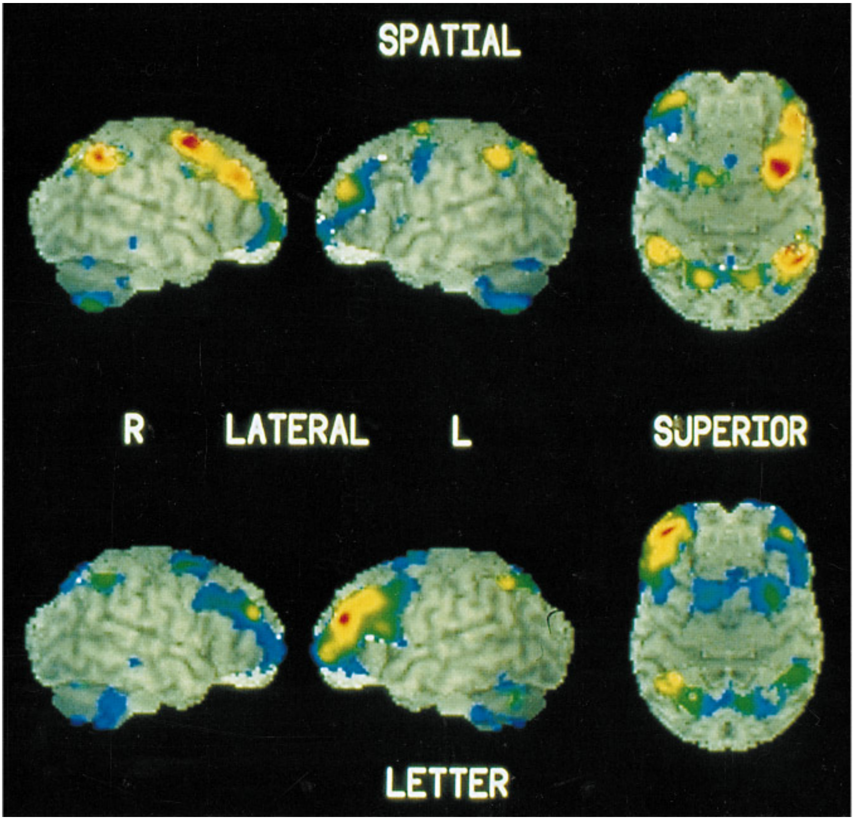

FIG. 5. PET activations for the spatial and verbal memory conditions (two-dimensional experiment), each with the control activations subtracted out. Shown in the figure are left and right lateral views as well as a superior view. The PET activations, shown in color, are superimposed on a surface rendering of a brain created from a standard MRI image (not one of the subjects in the experiment). The rainbow color scale used to represent activations ranges from blue to red (highest level of activation). The color scale directly reflects the significance of the activation, with *t* values ranging from 1.65 to 7.00, and values above 7.00 displayed as the peak red color.

tional regions are in the dorsolateral prefrontal cortex, areas routinely activated when computations must be performed on working memory.

A regional level of analysis of the spatial task yields a similar picture. The task activated a few familiar, distinctly spatial regions, including right-hemisphere premotor cortex and two sites in the right-hemisphere posterior parietal cortex. The task also activated regions not found in the initial spatial task, including homologous regions in left-hemisphere posterior, parietal cortex (again revealing some bilaterality), and dorsolateral prefrontal cortex (again reflecting computations performed on memory). Interestingly, in the spatial task most of the prefrontal activation was in the right hemisphere, whereas in the verbal task virtually all of the prefrontal activation was in the

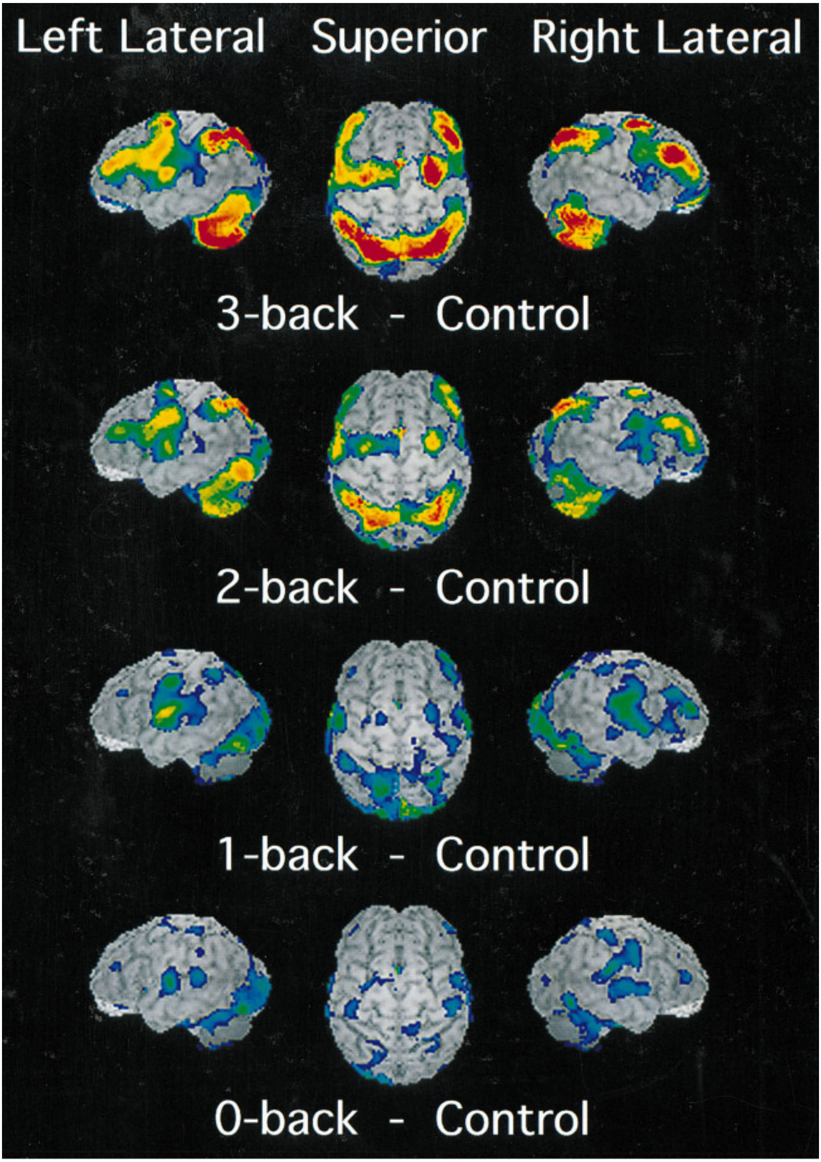

FIG. 10. PET activations for the four memory tasks in the 0123-back experiment, each with the control condition activation subtracted. Shown in the figure are left and right lateral views as well as a superior view. The PET activations, shown in color, are superimposed on a surface rendering of a brain created from a standard MRI image (not of one of the subjects in the experiment). The rainbow color scale used to represent activations ranges from blue to red (highest level of activation). The color scale directly reflects the significance of the activation, with t values ranging from 1.65 to 7.00, with values above 7.00 displayed at the peak red color. (Reproduced from Jonides *et al.*, in press.)

left hemisphere. The same kind of asymmetry holds for the posterior parietal area. Thus, in critical regions, lateralization is still present, with spatial mechanisms recruiting more right-hemisphere areas, and verbal mechanisms recruiting more left-hemisphere regions.[3,4]

Behavioral findings. Both two-dimensional experiments produced behavioral results in line with the subtraction logic. In the spatial versus object study, the difference in accuracy between the control and memory conditions was 88% vs 78% for the spatial task, and 84% vs 77% for the object task. In the spatial versus verbal experiment, the difference in accuracy between the control and memory conditions was 95% vs 90% for the spatial task, and 97% vs 86% for the verbal task. Whether distractors in the spatial task were near or far from targets was varied only in the spatial versus object experiment, and there accuracy was 98% for far distractors compared to 61% for near distractors. Again, our behavioral results manifest most of the desired patterns of results for subtraction-based neuroimaging studies.

Some implications. The two experiments discussed in this subsection make three general points.

1. They provide additional dissociations (quantitative rather than qualitative in one case) between spatial and object working memory and between spatial and verbal working memory.
2. They provide converging evidence that certain regions are part of the network for the various kinds of working memories. To focus on the strongest cases, two different studies show that: the right-hemisphere posterior parietal and premotor regions are part of the network for spatial working memory, whereas the left-hemisphere posterior parietal and Broca's regions are part of the network for verbal working memory. In that Broca's area is adjacent to the left-hemisphere premotor cortex, there is substantial symmetry between the regions involved in verbal and spatial working memory.
3. The dorsolateral prefrontal cortex may be critically involved when computations must be performed on the contents of working memory.

These three points have implications for cognitive analyses. Points (1) and (2) provide further support for distinguishing between verbal and spatial

---

[3] Some areas found in the initial verbal and spatial tasks were not significantly activated in the current experiment. Such ''missing'' areas include the left-hemisphere premotor and supplementary motor regions and the right-hemisphere cerebellum for the verbal task, and right-hemisphere premotor and occipital regions for the spatial task. The missing areas for the verbal task may all mediate rehearsal, and they may have been subtracted out because the verbal-control condition unintentionally involved some rehearsal. That is, even though the three target items were constant in the control condition, some subjects reported trying to keep them active. A similar story may hold for the missing premotor region in the spatial task.

[4] We have not performed a comparable two-dimensional study contrasting verbal and object working memory for the following reason. If the same stimuli contain object and verbal information, then the object information must be relatively verbalizable, and some subjects may choose to code this information verbally.

working memory at the cognitive level. Point (2) also suggests the need to distinguish between storage and rehearsal mechanisms (a point we develop in the next section). Point (3) suggests a qualitative difference between the processes that sustain basic working memory and the processes involved when computations are performed on working memory. This supports the distinction in Baddeley's (1986) model between the buffers on the one hand and the central executive on the other, and is contrary to proposals that combine the two functions into a pool of common resources (e.g., Carpenter, Just, & Shell, 1990). However, as argued elsewhere (Jonides, 1995), the ''central executive'' seems to include a set of distinct higher-level processes, not a single factor.[5]

THE ARCHITECTURE OF WORKING MEMORY: STORAGE VERSUS REHEARSAL

Different Circuits for Storage and Rehearsal of Verbal Working Memory

We have already mentioned that the left-hemisphere regions activated in our verbal working-memory tasks can readily be divided into two groups: a posterior parietal region that presumably mediates verbal storage, and a set of anterior speech regions (Broca's area, premotor, and supplementary motor areas) that presumably mediate verbal rehearsal. The argument for this division is based on the known functionality of these regions. Patients with lesions in the posterior parietal region have a selective deficit in remembering verbal information for brief periods of time (e.g., Shallice, 1988), whereas patients with lesions in the anterior speech regions (particularly Broca's area) have articulatory difficulties in explicit speech (e.g., McCarthy & Warrington, 1990).

Another PET study was explicity designed to provide converging evidence for the distinction between storage and rehearsal. The experiment included the three conditions presented in Fig. 6. The memory condition was like that in the verbal condition of the previous study, except that it is a ''2-back'' task—subjects decided whether the current letter was identical to the one 2-back in the series. There were two control conditions, both of which involved the same sequence of events as in the memory condition. In the search control, subjects decide whether each letter matched a single target letter identified at the beginning of the experiment. When this control was subtracted from the memory condition, we expected to find the predominantly left-hemisphere

---

[5] The experimental contrast supporting Point (3) involves not only different tasks, but different subjects and different PET scanners as well. In a more recent experiment, the same subjects were scanned in the same scanner while performing two verbal working memory tasks: an item-recognition task similar to the one used in the first study we described, and a 2-back task (i.e., ''Respond affirmatively only if the current letter matches the one two back in the series''). Again we found activation in the dorsolateral prefrontal area in the back task but not in the item-recognition task (Reuter-Lorenz, Jonides, Smith, Hartley, Cianciolo, Awh, Marshuetz, & Koeppe, 1996).

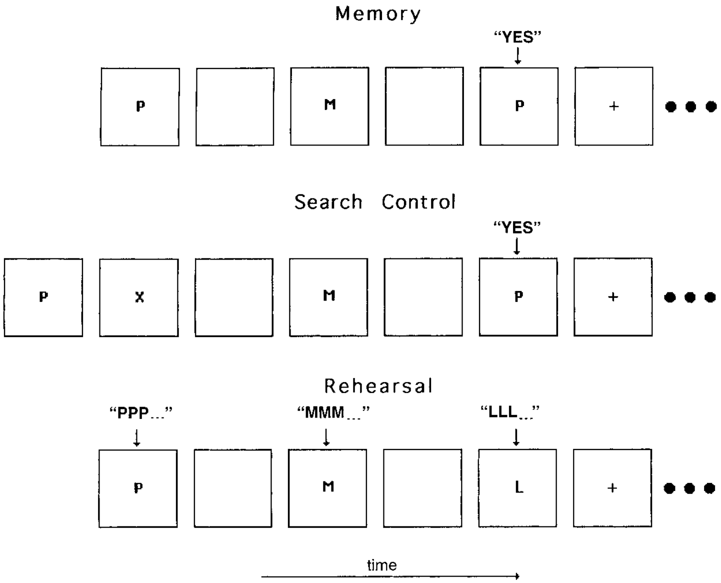

FIG. 6. Schematic representations of trials in the 2-back memory task (top panel), search control (middle panel), and the rehearsal control (bottom panel). (Adapted from Awh *et al.,* 1996.)

regions that were active in our previous verbal studies, including Broca's area, the premotor area, and the supplementary motor area. But the main purpose of this study hinged on the second, or rehearsal, control: subjects silently rehearsed each letter presented until the next one appeared, and then rehearsed that one. Subtracting this control from the memory condition should have removed much of the rehearsal circuit. (See Awh *et al.*, 1996, for details.)

The results were largely as expected. Consider first the results of subtracting the search control from the memory condition, which are presented schematically at the top of Fig. 7. In line with our previous studies of verbal working memory, most of the activated areas are in the left hemisphere, though there are some in the right hemisphere as well. More specifically, the main left-hemisphere areas that were significant in our initial item-recognition task are again significant here also, including the posterior parietal area, and the trio of speech regions that presumably mediate rehearsal—Broca's, premotor, and the supplementary motor area.

The subtraction of particular interest is between the memory condition and rehearsal control, and it is presented in the bottom of Fig. 7. As expected, a substantial part of the left-hemisphere rehearsal circuit has been lost. Neither Broca's area nor the premotor area is significantly active. However, the supplementary motor area continues to be active. That we were unable to subtract out the entire rehearsal circuit may be due to the fact that our rehearsal condition was not sufficiently demanding. Under the rehearsal control, sub-

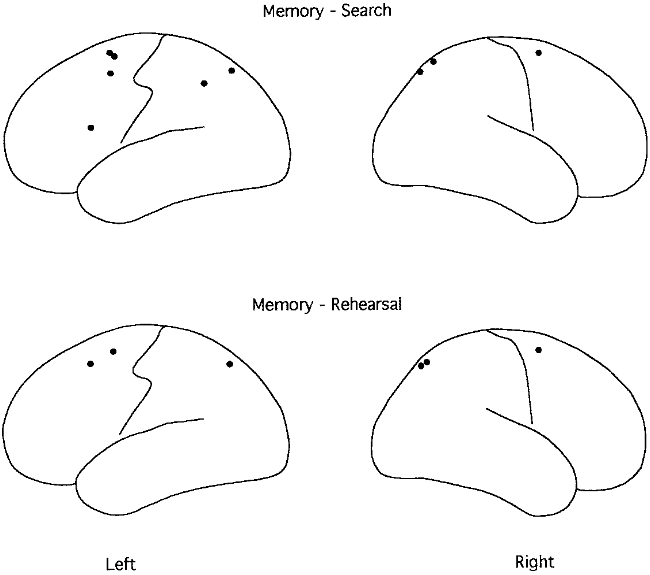

FIG. 7. Schematic representations of PET activations in left and right hemispheres, for the storage versus rehearsal experiment. The cortical activations for the memory-minus-search subtraction are in the top of the figure, while the cortical activations for the memory-minus-rehearsal subtraction are in the bottom of the figure. Several activations are not shown because they were subcortical; in the memory-minus-search subtraction, these subcortical activations were in the left- and right-hemisphere cerebellum and the right-hemisphere cerebellar vermis; in the memory-minus-rehearsal subtraction, the subcortical activations were in the right-hemisphere cerebellum, cerebellar vermis, and thalamus. Also not shown in the schematic is the anterior cingulate, which was activated in the memory-minus-rehearsal subtraction.

jects rehearsed one letter at a time, whereas under the memory condition they rehearsed two at a time; had the rehearsal requirements been identical, perhaps our subtraction would have eliminated the rehearsal circuit entirely.[6]

By and large, these results confirm our division of the relevant left-hemisphere areas into a verbal (phonological) storage buffer in the posterior portion of the brain and a verbal (phonological) rehearsal system in the anterior portion. An experiment by Paulesu, Frith, and Frackowiak (1993) provides additional support for this division. They contrasted an item-recognition task and a rhyming task, where the latter presumably involves some of the same

---

[6] In the memory-minus-search subtraction, there were areas of activation in addition to those found in the initial verbal experiment. These additional areas include a left-hemisphere cerebellar site and four right-hemisphere regions: among the latter are two regions in posterior parietal, one in the premotor area, and one in the cerebellum. The memory-minus-rehearsal subtraction also revealed some areas that were not found in our initial verbal experiment. These additional areas include five right-hemisphere regions: two in posterior parietal, one in the supplementary motor region, one in the thalamus, and one in the cerebellum.

phonological processes that are used in verbal rehearsal. Subtracting the rhyming from the memory task left the left-hemisphere posterior parietal cortex activated but eliminated the activation in Broca's area. It seems, then, that the brain honors the buffer vs rehearsal distinction so common in cognitive models.

Finally, we note that the behavioral data in this experiment were again in line with subtraction logic. Accuracy in the search-control condition was 100%, while that in the memory condition was 96%. Obviously, the true difference between the two conditions is being obscured by a ceiling effect. The latencies provide a clearer picture, as subjects averaged 552 ms in the search-control condition versus 662 ms in the memory condition.

Different Circuits for Storage and Rehearsal in Spatial Working Memory

While the notion of verbal rehearsal is relatively straightforward, that of spatial rehearsal is less clear. Our hypothesis is that rehearsing a spatial location in working memory involves (in part) selectively attending to a representation of that location (e.g., Awh *et al.*, 1995). We have not yet completed a neuroimaging experiment that evaluates this hypothesis, but there are existant neuroimaging data that can be used to test it partially.

The test we have in mind is a comparison between the imaging results obtained in our studies of spatial working memory with those obtained in studies of selective attention. Our studies of spatial working memory routinely activate right-hemisphere areas in posterior parietal and frontal cortices. The frontal region most consistently activated is the premotor region. The posterior parietal regions activated include two distinct sites, an inferior one and a superior one (we have not bothered with this distinction earlier, but we need it now). It turns out that two of the regions just mentioned—the premotor one and the superior, posterior parietal one—are also routinely activated in studies of selective attention.

We can illustrate the latter point with a PET experiment by Corbetta, Miezin, Shulman, and Petersen (1993). They contrasted an attention task with a control condition. In the attention task, subjects had to fixate on a central point while detecting targets presented in the periphery to the left or right of fixation. The vast majority of the targets appeared in a predictable location, which meant subjects could shift their attention toward the location of the upcoming target. (Independent behavioral results indicated that subjects do indeed shift their attention in these circumstances). Under a control condition, subjects again fixated on a central point, but now simply had to detect a target presented centrally and ignore randomly presented peripheral targets. Subtracting the control from the attention condition resulted in four significant areas of activation. Two of these correspond to the spatial-memory regions we emphasized above—the premotor region and superior, posterior parietal region. One of the other regions obtained by Corbetta *et al.* (1993) is in the anterior cingulate, which often is activated in our spatial-memory studies (though not always significantly so).

All in all, there is a promising correspondence between Corbetta *et al.*'s (1993) activations and the spatial-memory activations we presented earlier. Moreover, a recent review of all relevant neuroimaging experiments provides robust evidence for corresponding areas being activated in spatial-attention and spatial-memory tasks (Awh & Jonides, 1996). Interestingly, though, one area that is consistently activated in spatial-memory but not spatial-attention tasks is the inferior, posterior parietal cortex (right hemisphere) that we have mentioned repeatedly. This area might well serve as a spatial buffer (as we have assumed throughout), while the superior, posterior parietal site plus the premotor one may form part of a network for spatial rehearsal. While we need a neuroimaging experiment to test the above detailed claims, a recent behavioral study supports the basic idea that rehearsing a spatial location involves selectively attending to it. In this experiment (Awh, 1996), subjects performed two tasks simultaneously. One task required spatial working memory, as subjects had to remember the position of a single letter during a 500-ms retention interval. The other task required subjects to make a discriminative response to a non-letter target that was presented during the retention interval of the spatial-memory task. The major variation was whether the target in the discrimination task appeared in the same location that subjects were presumably rehearsing in the memory task. Subjects responded faster to the discrimination target when it appeared in the rehearsal position, supporting the notions that subjects were selectively attending to this position, and that selective attention is part of spatial rehearsal.

An Updated Architecture of Working Memory

General architecture. We can summarize the results reported in this section and the preceding one by the flow diagram in Fig. 8. It contains some basic cognitive components of working memory along with an indication of how they are implemented in the brain.

We know from other work that any visual input can be coded with respect to both its spatial and its object content (the former is implemented by the occipital–parietal pathways that comprise the ''where'' system, and the latter by the occipital–temporal pathways that make up the ''what'' system—e.g., Mishkin *et al.*, 1983). We also know that certain visual inputs, those that are linguistic symbols, can be translated from a visual representation into a phonological one. We therefore have three kinds of information—verbal (phonological), spatial, and object—that can serve as inputs to working memory. Three different working-memory systems appear to handle the three kinds of inputs. The verbal and spatial systems each appear to include a buffer and rehearsal process. The buffers for the verbal and spatial systems are housed (in part) in posterior parietal regions, with the verbal buffer being left lateralized and the spatial buffer being right lateralized (see the bottom of Fig. 8). The rehearsal process for the verbal system is mediated by a circuit in the left-hemisphere frontal cortex that includes the speech regions (e.g., Broca's area), whereas the rehearsal process for the spatial system may be

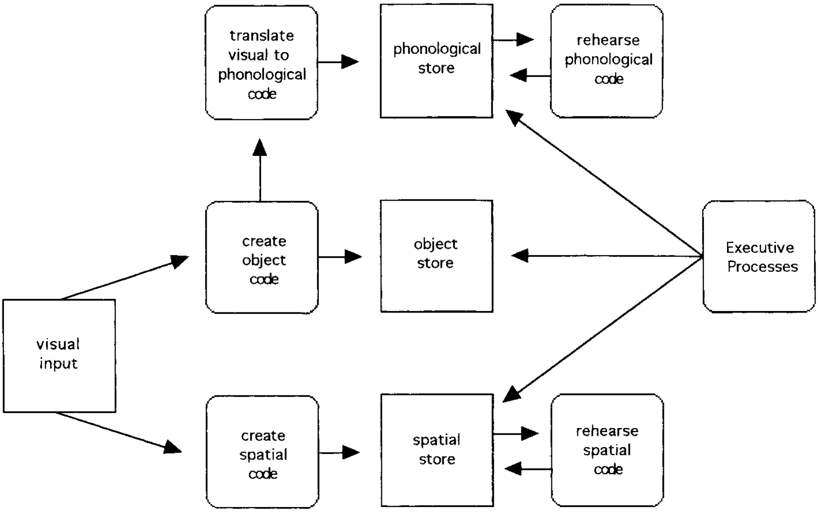

Phonological Store: Predominantly Posterior Parietal (left lateralized)
Phonological Rehearsal: Broca's area, premotor and supplementary motor area (left lateralized)

Spatial Store: Predominantly Posterior parietal (right lateralized)
Spatial Rehearsal: Premotor area (right lateralized)

Executive Process: Dorsolateral prefrontal cortex

FIG. 8. A flow diagram containing some of the basic cognitive components of working memory. The information at the bottom of the figure indicates where in the brain some of these components are implemented.

implemented by a parietal–frontal circuit that includes the premotor area (see bottom of Fig. 8). For the object system, as yet we have no evidence for a storage-rehearsal distinction. We know only that two left-hemisphere regions may be involved, the inferotemporal and posterior parietal regions.

Depicted at the right of Fig. 8 are a set of ‘‘executive’’ processes that can operate on any of the working-memory buffers. These executive processes include contextual coding of the contents of working memory, such as the temporal coding required in our 2- and 3-back tasks. These executive processes are implemented, in part, by regions of the dorsolateral prefrontal cortex. Although we assume that the same strategic processes can apply to different kinds of working memories, it may turn out that some strategic processes are material specific, too.

Two aspects of this model are worth comment. First, the model suggests that the input entering verbal working memory is strictly phonological, even though the initial representation was visual. This in turn suggests that here should be no difference in the activations underlying verbal memory for information presented visually versus auditorally, once input differences have been subtracted out. We have tested this suggestion (Schumacher, Lauber, Awh, Jonides, Smith, &

Koeppe, 1996). Using the 3-back task and control conditions described earlier (see the discussion on two-dimensional tasks), we varied whether the input modality was visual or auditory. Once the controls were subtracted from the memory conditions (which eliminates activations due to input modality), there was virtually no difference between the activation patterns for the two memory conditions. Thus, this aspect of the model seems tenable.

The other aspect of the model that deserves comment is the distinction between perceptual and memorial representations. In line with most cognitive models, we distinguish between an initial spatial (object) representation that is created by the perceptual system, and a subsequent spatial (object) representation that is maintained in a short-term buffer (see Fig. 8). However, at this point in time, it is possible to argue that, for spatial and object information, perceptual and memorial representations are implemented by the same neural regions. Thus, we have continually noted that the right-hemisphere posterior parietal region implements a spatial buffer, but this same region has also been shown to be involved in spatial perception; similarly, we have suggested that the left-hemisphere inferotemporal region is involved in storing object information, but this region may also play a role in object perception (see Smith *et al.,* 1995). These findings, however, are not based on within-experiment comparisons and are not very precise. Should subsequent work show that perceptual and (short-term) memorial representations are mediated by the same regions, we will want to change the model in Fig. 8 so as to eliminate the cognitive distinction between perceptual and memorial representations.

More detailed brain-cognition mappings and cognitive analyses. Our mappings between brain regions and cognitive processes are at a gross level: an area of cortex containing between roughly 25,000 and 1,000,000 neurons is mapped onto a very general cognitive function like a spatial buffer. We are obviously a long way from understanding the level of neural computation, or even from the level of computational analysis that charaterizes much of cognitive psychology. However, we can take a step toward more precise mappings and analyses by considering single-cell research with nonhuman primates that has investigated brain regions analogous to those we found activated, and that has determined the computations carried out by some of the neurons in these regions. The relevant research includes single-cell recording studies of spatial and object working memory in monkeys, and the areas of particular interest include posterior parietal cortex, prefrontal cortex, and inferotemporal cortex.

Studies of spatial working memory have required animals to remember a particular spatial location for a few seconds. Recordings have been taken from posterior parietal cortex—active in all our spatial tasks—and prefrontal cortex—active in our two-dimensional spatial task (e.g., Funahashi, Bruce, & Goldman-Rakic, 1989; Goldman-Rakic, 1987). One finding of interest is that cells have been found that fire only when a location is being held in storage (i.e., during the retention interval), not when the location is presented. Furthermore, particular cells fired at high rates only when particular locations were being stored, and at reduced rates when adjacent locations were being maintained. These cells have been interpreted as broadly tuned location-memory units, and activity in these units may be part of the neural representation of working-memory storage or rehearsal for spatial location.

Comparable single-cell studies have been done on object working memory. In several studies (e.g., Desimone, Albright, Gross, & Bruce, 1984; Fuster & Jervey, 1981), the neurons of interest are located in inferotemporal cortex a region that we found active in our object working-memory task. In other studies, the relevant neurons are found in a part of prefrontal cortex that connects directly to the inferotemporal region. In all cases, the neurons fire only during a retention interval, and at least in the inferotemporal case, the neurons appear to fire only to particular visual features, such as a part of an object or a feature of the entire object (Desimone *et al.,* 1984). These cells may be viewed as object-memory units, and activity in these units may be part of the neural representation of working-memory storage or rehearsal for objects.

Obviously, we cannot rely on nonhuman studies to elucidate verbal working memory. But by analogy with the preceding research, there may be neurons in some of the frontal speech regions we have found active that respond only to certain phonological features and that may function as part of the neural representation of verbal working memory.

The preceding points suggest a way of relating neural and cognitive computations. At the cognitive level, something that is preserved as a unitary item is decomposed into constituent features (the pronunciation of a letter may be composed of phonological features; an object may be composed of shape, color, and texture features; and a spatial location may be a mixture of certain canonical locations.) Each such feature is a cognitive entity, and such entities are already part of many existent cognitive models. Also, each such feature may be detected by specialized neurons in the appropriate cortical region (the way shape and color features are detected by neurons in the inferotemporal region), and this information may then be passed to other neurons that are capable of briefly sustaining their activity in the absence of external stimulation and hence can serve a storage/retrieval function. In this way, one can develop models of working memory that offer computational accounts at the cognitive level, along with proposals for how to implement these computations neurally, where these proposals are compatible with human imaging data and nonhuman single-cell data.

CONTINUOUS VARIATIONS IN BASIC WORKING-MEMORY COMPONENTS

Logic of Parametric Variation

Thus far, every result we have presented hinges on the subtraction method. This method, introduced by Donders in 1868, has produced many important behavioral results (see, e.g., Posner, 1978), but it rests on a problematic assumption. The assumption at issue is called ‘‘pure insertion’’ (Sternberg, 1969): the claim that the addition of a particular processing stage to some task does not affect the operation of other stages in that task. For example, the difference between the memory and search-control tasks in our ‘‘*n*-back’’ studies is one of working memory, and the assumption of pure insertion says that adding this memory requirement has no effect on processing stages concerned with perception and response. This assumption has been challenged in behavioral studies (dating back to Kulpe, 1895), and it may sometimes be misleading in neuroimaging experiments as well. In our *n*-back tasks, for example, adding a memory requirement may also add to the perceptual analysis needed, and consequently the activations present in a memory-minus-control subtraction could reflect perceptual networks in addition to mnemonic ones.[7]

To avoid these potential problems, we can augment the subtraction approach with a method of ‘‘parametric variation.’’ Here, one varies an experimental factor that presumably affects the operation of a single processing stage (e.g., the number of items to be remembered in an item-recognition task), and determines if this factor affects regional activation in a systematic way. This approach to neuroimaging has something in common with Sternberg’s (1969) additive-factors approach to reaction-time measurements, and it has already been used successfully in neuroimaging studies of long-term memory (e.g., Grasby, Frith, Friston, Bench, Frackowiak, & Dolan, 1993). We employed it in the experiments described in this section.

Effects of Continuous Variations in Verbal Working-Memory Load

An n-back experiment. A natural for parametric variation is our *n*-back task. We can vary the value of *n*, thereby varying memory load, and examine the effects on activations in different neural areas. Two hypotheses about possible activation-effects are of interest:

1. Continuous increases in memory load lead to continuous increases in activation in a fixed number of memory areas (these areas being the ones that were activated in our previous studies of verbal working memory); and
2. Continuous increases in memory load lead to new areas being recruited for the task.

The ‘‘more processing in fixed areas’’ hypothesis is compatible with the claims that: (a) previous research has identified some of the basic neural components of verbal working memory, and (b) activation differences in different memory tasks arise because these basic components are differentially utilized in the tasks. These claims about neural processes mesh well with the

---

[7] Raichle (1996) points out that large-scale violations of the subtraction logic—as when the addition of a processing requirement leads to a strategy change—are easy to spot in neuroimaging experiments because they produce substantial, qualitative changes in activation patterns. However, a violation like the increase-memory-recruits-increased-perception one just noted in the text is harder to detect, and hence more pernicious.

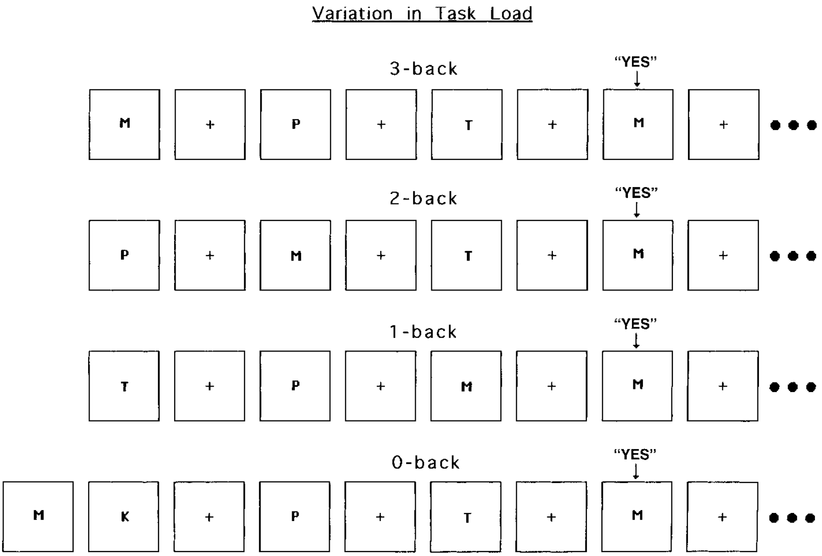

FIG. 9. Schematic representations of trials in the four memory tasks of the 0123-back experiment. The tasks differ in working memory load, varying from 0-back to 3-back. In addition to these conditions, the experiment also included a control condition (not shown), in which letters were presented in sequence and subjects simply had to respond with a keypress upon presentation of each letter. (Adapted from Jonides *et al.*, in press.)

cognitive notion of a small, fixed number of architectural components (as presented at the end of the previous section). The second hypothesis—the ''more areas'' hypothesis—does not require that there be a basic set of neural mechanisms, and it seems to map less well to the cognitive notion of a small, fixed number or architectural components.

To test these hypotheses, we performed a PET *n*-back experiment that used four different levels of *n*: 0, 1, 2, and 3. The four tasks are presented in Fig. 9, and they are similar to the *n*-back tasks discussed earlier. In the most difficult task, subjects had to match each letter to the one 3-back in the series. In easier tasks, subjects had to engage in 2-back and 1-back matching. And in the easiest task, subjects matched each letter against a constant target letter that was specified at the outset (this is 0-back). In addition to these four memory conditions, we included a control condition in which subjects simply pressed a response button when each letter appeared. This control was used as a baseline whose PET activations were subtracted from those of the memory conditions. Note that subtraction is being used here not to isolate processing stages at either the neural or cognitive level, but rather to correct for individual differences in physiological responsivity (for details, see Jonides, Schumacher, Smith, Lauber, Awh, Minoshima, & Koeppe, in press).

Figure 10 presents the results as PET activations superimposed on a surface rendering of a brain, separately for each of the four tasks. There are many

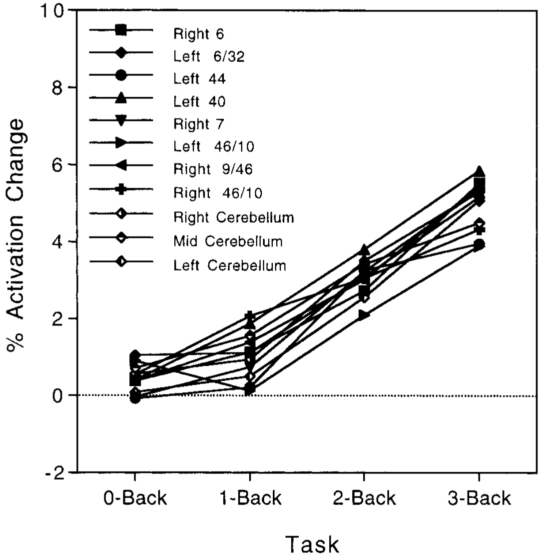

FIG. 11. Percentage change in activation in the 0123-back study as a function of memory load for eleven regions of interest. The regions are ones that previous studies (Awh et al., 1996; Schumacher *et al.,* 1996) had found active in 2-back and 3-back verbal memory tasks. Values are the mean change in activation across spheres of 5.4-mm radius. (Adapted from Jonides *et al.,* in press.)

more areas of reliable activation in the 2- and 3-back tasks than in the 0- and 1-back tasks. Statistical analysis bears this out: there are 22 significant sites of activation in both 2-back and 3-back, but only 2 significant areas in 0-back and 3 significant ones in 1-back. At face value, then, these results support the more-areas hypothesis: as memory load is increased, more areas are recruited to do the task.

Further analyses, however, indicate that this is *not* the case. Rather, what is happening is the following: with increasing memory load, regions that *were active but below the threshold of statistical significance* are more likely to pass that threshold. Support for this claim is given by the analysis presented in Fig. 11. There we have plotted the change in activation (bloodflow) as a function of memory load, separately for 11 different areas of interest. These regions are our estimate of the basic verbal-working-memory areas: they were selected on the basis of being significantly active in the *n*-back experiments discussed in the previous sections. The most striking aspect about the functions in Fig. 11 is that nearly all of them show monotonic increases in activation with memory load. (Statistical analysis shows that every single function manifests a significant linear trend.) These results support the hypothesis of more processing in fixed areas.

To provide more definitive support for the fixed-areas hypothesis, we need to show that areas *not* involved in working memory do *not* show monotonic

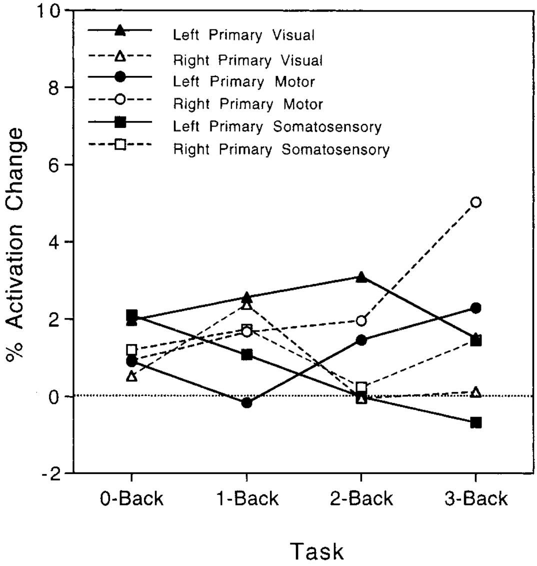

FIG. 12. Percentage change in activation in the 0123-back study as a function of memory load for six non-memory regions. The regions include areas in primary-visual, primary-motor, and primary-sensorimotor cortex. (Reproduced from Jonides *et al.*, in press.)

increases in activation. To demonstrate this, first we selected six areas that *a priori* should not be involved in working memory—two areas in primary visual cortex, two in primary motor cortex, and two in primary somatosensory cortex. Then we determined the activation in these areas as a function of memory load. The results are displayed in Fig. 12. Not one of the areas shows a significant effect of memory load, let alone a monotonic increase. Thus monotonic increases in activation occur only for working-memory areas, as assumed by the more-processing-in-fixed-areas hypothesis.[8]

Though this hypothesis provides a good overall account of the data, there are two cases in which we see something like the recruitment of a new area. One of the cases involves the dorsolateral prefrontal cortex. We noted earlier that this region may mediate executive processes, and we suggested that it underlies the temporal coding needed in *n*-back tasks. However, an intuitive analysis of the task suggests that temporal coding may not be needed in the 0-back and 1-back versions (subjects may simply match the current letter to the constant target in 0-back, or to the most activated item in 1-back); consequently, the dorsolateral prefrontal cortex may not be much activated until the 2-back task. We found support for this line of reasoning when we

---

[8] In Fig. 12, there appears to be a modest activation increase with memory load for one of the motor regions. But this apparent increase is likely due to the fact that the specific motor region we selected is adjacent to the premotor region, which is one of our working-memory areas.

examined the activation changes in dorsolateral prefrontal cortex: for 0-back, 1-back, 2-back, and 3-back, the activation changes are 1.39, 1.34, 3.44, and 5.55%. These changes are characterized by a step function, which fits with the idea that the activation is mediating a new cognitive function, which in turn fits with the idea that a new area is being recruited when a new task demand is added. A step function also characterized the activation changes in Broca's area: for 0-back, 1-back, 2-back, and 3-back, the activation changes were 0.97, 1.07, 4.40, and 5.00%. Since Broca's area is part of the rehearsal circuit, presumably there is little need for rehearsal in 0-back and 1-back tasks (a suggestion that also fits with intuition).

In sum, this experiment makes three major points:

1. The method of parametric variation produces results that converge with the subtraction method. The basic working-memory areas identified in our previous subtraction-based studies are the areas that show monotonic increases in activation in the current parametric-variation experiment;[9]
2. By and large, continuous increases in memory load lead to continuous increases in activation in a fixed set of areas; and
3. There is an important amendment to the preceding point, in that an increase in memory load may sometimes be accompanied by a new task requirement, which can lead to a marked increase in the activation of an area.

The kind of cognitive proposal that fits best with these points is one that assumes a fixed cognitive architecture with a relatively small number of basic components, and that explains variations in behavioral performance in terms of variations in the activity of the basic components.

Behavioral findings. We are again interested in whether our behavioral data manifest certain kinds of findings including:

1. Accuracy is relatively high under all conditions.
2. Performance declines continuously with continuous increases in task load.
3. There are correlations between behavioral performance and activation measures.

The desired accuracy and correlational results are the same as those considered with the subtraction-based studies. The second type of finding is an extension of the contrast between target and control conditions used in subtraction-based studies to test the subtraction logic. Now, not only must performance in 1-back be less than in 0-back, but performance in 2-back must be less than in 1-back, and so on. This set of inequalities tests the logic of

---

[9] There is, however, one important discrepancy between the current results and our earlier ones. While our previous studies of verbal working memory show a left-hemisphere lateralization (sometimes qualitatively, sometimes only quantitatively), the present results reveal completely bilateral activations (see Fig. 10). We are unsure of the source of this discrepancy, but it is worth noting that among the previous studies, those with discrete trials (the item recognition task) generally showed more left lateralization than those using continuous presentation (the back tasks).

parametric variation, and it is sufficiently diagnostic to serve also as a behavioral signature of the memory processes of interest.

The behavioral findings from the preceding experiment stack up well against the above criteria. Accuracy was in the range of 88 to 96% correct, clearly indicating that subjects were executing the processes of interest. More impressively, both accuracy and speed declined monotonically with increases in task load.

The correlational data from this PET experiment were not conclusive about relationships between activation and performance, largely because of high variability in activations, small number of points per subject, and low variation in performance measures. However, we have collected correlational data from a related study that used the same 0-,1-,2-, and 3-back tasks, but the neuroimaging technique was functional magnetic resonance imaging (fMRI), not PET. Because experiments can be conducted using fMRI without a contrast agent, one can collect numerous observations on an individual and look for performance-activation correlations within each subject. Two frontal areas were scanned, Broca's area and dorsolateral prefrontal cortex, and within these regions we determined the correlations between mean reaction time and mean fMRI signal within a block of trials. The correlations were on the order of .50, clearly significant. These correlations reflect, in part, the fact that both reaction time and activation increased monotonically with increases in memory load. (For details, see, Braver, Cohen, Jonides, Smith, & Noll, 1997).

In sum, the behavioral findings are in line with the claim that each increase in memory load leads to an increase in the intensity and/or duration of mnemonic processes.

Determining the Temporal Course of Memory-Based Activations

It would be useful to determine the temporal course of the activations we observed in the previous parametric study, for at least two reasons:

1. In the previous study, we interpreted the continuous increases in activation with increasing load to mean that the area is processing more intensely. An alternative interpretation is that the area is simply working longer.
2. Perhaps more importantly, the temporal course of an activation tells us something about that area's function. Consider an extreme example. In an *n*-back task, if we find that an area is active when an item is presented, but not during the interstimulus interval, we would not want to attribute a storage function to that area.

Given the current state of technology, we cannot use PET to address these two issues. Since PET accumulates radioactivity counts over a 60-s period, we cannot tell more-intense from longer-duration processing, nor can we discriminate between events that occur at different points during the recording interval.

It is possible, however, to address these two issues with fMRI. The latter technique now affords a temporal resolution of about 2.5 s. Furthermore, images

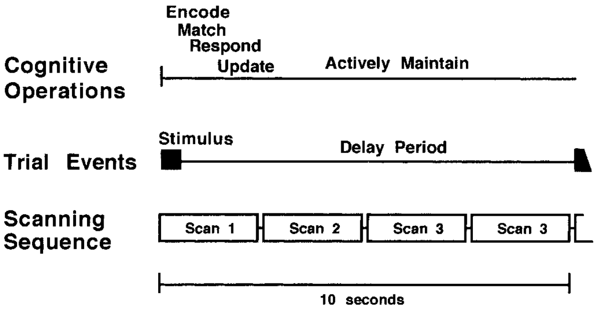

FIG. 13. Schematic representation of the sequence of events in the temporal dynamics fMRI study, and the relation of these events to the time course of the cognitive processes involved. The sequence of trial events is given in the middle of the figure. The bottom of the figure shows when the four scans occurred in relation to the trial events, while the top suggests when the various cognitive processes take place. (Reproduced from Cohen *et al.*, 1996.)

taken every 2.5 s can be time-locked to a stimulus event. These aspects allowed us to do the following variant of the preceding parametric study. We used the 0-, 1-, 2-, and 3-back tasks, but we increased the interstimulus interval to 10 s. This enabled us to scan subjects for four time periods during the interstimulus interval: 0–2.5 s, 2.5–5.0 s, 5.0–7.5 s, and 7.5–10 s. (For details, see Cohen, Perlstein, Braver, Nystrom, Noll, Jonides, & Smith, in press).[10]

Figure 13 gives some notion of how we can relate temporal patterns of activation to supposed functions of the activated areas. The basic idea is that different cognitive functions (or processes), have different temporal courses. The *n*-back task requires at least the following processes: (a) encoding a stimulus, (b) matching it to a representation of the item *n*-back, (c) responding on the basis of this comparison, (d) updating the *n* items that need to be kept in an active state and their temporal codes, (e) storing these items and their temporal codes, and (f) rehearsing these items and codes. We expected that:

1. Areas that mediate encoding and response (visual and motor areas) may show no effect of task load, and display transient activation (i.e., activation that beings relatively soon after stimulus onset and declines substantially before the interval is completed).
2. Areas that presumably mediate storage (posterior parietal areas) and rehearsal (e.g., Broca’s area) should show an increasing effect of load (as in the previous study), and display activation that is sustained during the entire interval.

---

[10] With regard to the logic of fMRI measurements, suffice it to say that: (1) An increase in regional neural activity leads to an increase in blood flow to that region; (2) not all the increased oxygen brought by the blood is used by the region, leading to a local increase in oxygenated hemoglobin, and a decrease in the relative amount of deoxygenated hemoglobin; and (3) changes in deoxygenated hemoglobin have paramagnetic properties, which affect the MRI signal.

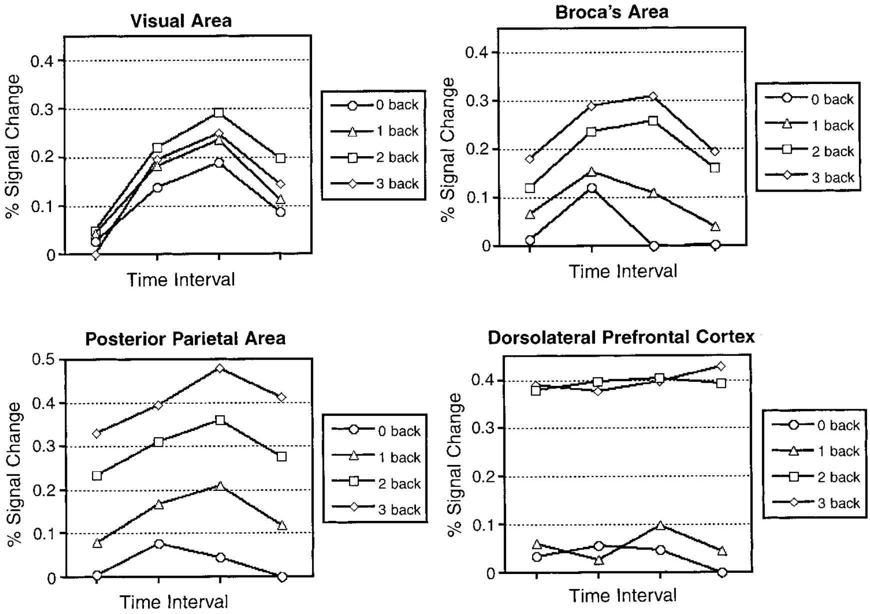

FIG. 14. Percentage of fMRI signal change in activation as a function of temporal interval, with working memory load as the parameter (temporal dynamics study). The results are presented separately for four different areas, which in descending order are: (1) an area in visual cortex, (2) a posterior parietal area, (3) Broca's area, and (4) an area in dorsolateral prefrontal cortex. (Adapted from Cohen *et al.,* 1996.)

3. Areas that presumably mediate temporal coding (dorsolateral prefrontal area) should show an increasing effect of load (as in the previous study), and display some decrease in activation as the interval lengthens.

Four sets of results are presented in Fig. 14. Each graph plots activation as a function of time period, separately for the four memory loads. The top left graph presents the results for a visual area that likely is involved in encoding the letter stimuli. Just as expected, there is no statistical effect of memory load, and the activation is clearly transient. (We found the same pattern when we examined activation in the primary motor areas.) Thus cognitive processes that need be performed only once during a trial—say, encoding and response—are associated with transient activation patterns at the neural level.[11]

The bottom left graph in Fig. 14 shows the results for a posterior parietal area that our previous studies led us to believe plays a storage function. As expected, activation increases continuously with memory load, and it is

---

[11] The hemodynamic response that drives the fMRI signal (see footnote 10) has a built in 4- to 5-s delay. Consequently the fMRI activations will not reach their peak until well into the interstimulus interval, even for areas that mediate perceptual processes.

relatively sustained through the interval. This is entirely consistent with the area playing the role of a storage buffer. But there is an important qualification. A posterior parietal area adjacent to that featured in Fig. 14 showed transient patterns of activation during the interval (though it manifested the expected effects of storage load). It is possible that the two different parietal regions have somewhat different functions.

The top right graph in Fig. 14 depicts the results for Broca's area, a prime candidate for mediating rehearsal. As expected, activation increases with memory load (in something of a step-function, thereby replicating our previous PET results). Somewhat surprisingly, the activations decrease as the interval lengthens. Perhaps the simplest interpretation of these results is that the area does indeed mediate rehearsal (there is considerable independent evidence for this point), but the rate of rehearsal decreases during a 10-s. interval (a suggestion that fits with intuition).

The bottom right graph in Fig. 14 is for the dorsolateral prefrontal cortex. Again, activation increases with memory load, with this increase taking a substantial step from 1- to 2-back (replicating in detail the previous PET results). This step-function fits with the idea that the area is involved in temporally coding the contents of working memory, since such codes may be needed only when the memory load is two or more items. What fits less well with this temporal-coding function, though, is the finding that all activations are sustained throughout the interval. We had thought that temporal coding would not consume the entire interval, but it is possible that such codes have to be maintained by a special process—different from the posterior process that maintains the actual contents of working memory—and that this maintenance is part of the job of dorsolateral prefrontal cortex (Cohen, Braver, & O'Reilly, in press).

Finally, we note that the behavioral data in this experiment manifested the same kinds of findings that we obtained in the previous, parametric study. Accuracy was in the range of 90 to 98% correct, indicating that subjects were performing the processes of interest. Furthermore, accuracy and speed generally decreased with increases in task load; for accuracies, the decrease was strictly monotonic; for latencies, the mean for 0-, 1-, 2-, and 3-back were 729, 912, 1036, and 1018 ms, respectively. The apparent latency reversal for 2- and 3-back was not significant. All things considered, the behavioral data show great sensitivity to the parametric variations in memory load, just as the activation data do.

To summarize, the results of this experiment:

1. Replicate in detail the effects of memory load on activation, including the findings that for certain areas the increase in activation is a step-function;
2. Suggest that the above increases in activation reflect more intensive processing, not just longer-duration processing;
3. Show substantial differences between the temporal pattern of areas that perform operations needed only once during a trial (e.g., perception and

response) versus those that perform operations that need to be maintained through the trial (e.g., storage of the contents of working memory, and temporal updating and storage of temporal codes);

4. Suggest that rehearsal is not maintained at the same rate during a long (10 s) interval;
5. Suggest that a special process (presumably an executive process) not only handles the updating of temporal codes but also the maintenance of these codes.

The last three points are the products of our temporal analysis, and we note that such an analysis goes beyond what is possible with strictly behavioral methods. With the latter, if one wants to discover what is happening at various points in, say, a 10-s. retention interval, the experimenter must somehow interrupt the subject during that interval. Not so with neuroimaging—one can scan at various points without in any way interrupting the subject. In this sense, neuroimaging provides a picture of what previously has been hidden from view.

SUMMARY AND FUTURE DIRECTIONS

Summary

Our major empirical findings are as follows.

Different neural circuits mediate verbal, spatial, and object working memory, with spatial memory being right lateralized and verbal and object memory typically being left lateralized. The difference in lateralization is likely quantitative rather than qualitative (and appears to be greater in discrete than continuous tasks). In addition to lateralization, some of the cortical regions involved in each kind of working memory seem to be distinctive. Only in spatial working memory is there activation in an occipital region and an inferior frontal region; only in object memory is an inferotemporal area active; and only in verbal memory is there activation in Broca's area.

Within the circuitry for verbal working memory, there is a separation between components that underlie storage and those that mediate rehearsal. Regions in the posterior parietal cortex (particularly in the left hemisphere) seem to be involved in storage, whereas frontal regions appear to mediate rehearsal. The latter include three left-hemisphere regions known to be involved in higher-level aspects of speech: Broca's area, the premotor area, and the supplementary motor area. There is some suggestive evidence that spatial working memory has a parallel storage–rehearsal structure, with storage again being mediated by posterior parietal regions (though now right rather than left lateralized), and rehearsal perhaps again involving the premotor area in frontal cortex (though right rather than left lateralized).

In contrast to storage and rehearsal functions, the processes that operate on the contents of working memory (executive processes) appear to be mediated by the dorsolateral prefrontal cortex. In particular, this region was activated when our subjects had to temporally code the contents of working memory rather than just store the materials.

All of the preceding results were initially obtained in paradigms that relied on the subtraction method. At least for verbal working memory, the same regions are activated when the method of parametric variation is employed. Furthermore, continuous increases in memory load lead to increases in activation in these verbal memory areas, but not in other regions.

Finally, in addition to characterizing an area by its degree of activation in a verbal working-memory task, we may also characterize it by its temporal pattern of activation. Areas that mediate processes that are needed only once during an interval (e.g., encoding and response processes) show transient activation during the interval; areas that mediate processes needed throughout the interval (e.g., storage of the verbal contents or their temporal codes) show sustained activation during the interval.

Many of these findings have direct consequences for cognitive models. For a cognitive model of working memory to be consistent with the neural findings, it must distinguish the three types of working memory; provide a storage versus rehearsal structure to each working memory; posit a separate set of executive processes; explain variations in specific tasks as a function of variations in the extent to which the above-mentioned architectural components are involved in the task; and posit different temporal dynamics for different kinds of cognitive processes.

Other Kinds of Working Memories

While we have focused on verbal, spatial, and object working memory (particularly the former two), there may be other kinds of working memories as well. For one thing, there may be a separate working memory for each major sensory system. Indeed, there already is neuroimaging evidence for a short-term buffer for auditory information (see, e.g., Zatorre, Halpern, Perry, Meyer, & Evans, 1996). Assuming that such sensory buffers exist, note that they differ from the verbal and spatial working memories on which we have focused. Sensory buffers, by definition, store information from only their own input modality, while we have already provided PET evidence that verbal working memory is the same for visual and auditory inputs (Schumacher *et al.*, 1996). And spatial working memory may turn out to be similarly amodal. Furthermore, it may even be a bit of a misnomer to refer to sensory buffers as ‘‘working’’ memories, since the only work they do is to aid in the processing of identity information.

Another potential kind of working memory is more abstract than verbal, spatial, and object working memory. While verbal working memory can be used to store intermediate products in many higher mental process—like mental arithmetic—its representations are phonological, and hence do not carry meaning. The kind of working memory needed to support critical operations in reasoning, problem solving, and language understanding must be meaning-based and capable of extracting the same meaning from verbal or spatial inputs. Such a working memory would be ‘‘propositional,’’ and its neural basis remains elusive (see Jonides & Smith, in press, for further discussion).

Executive Processes

We close with a comment about executive processes. In the research presented here, the only executive process we considered is temporal tagging, and we suggested that it involves the dorsolateral prefrontal cortex. Other neuroimaging experiments have studied a variety of executive processes, including: monitoring the contents of working memory (Petrides *et al.*, 1993a,b); switching between two tasks that both require working memory (e.g., D’Esposito, Detre, Alsop, Shin, Atlas, & Grossman, 1996; Rubenstein, Meyer, & Evans, 1994); applying a complex rule that may be kept active in working memory (Smith *et al.*, in press); and planning a set of moves in a problem-solving task (e.g., Owen, Doyan, Petrides, Evans, & Gjedde, 1994). All of these studies found activation in dorsolateral prefrontal cortex. How can such different cognitive functions be mediated by the same neural area?

Perhaps different parts of this prefrontal region subserve the distinct functions mentioned above, but current neuroimaging techniques lack the spatial resolution to distinguish these different regions. Alternatively, perhaps these functions have some component in common, and this common component is what activates dorsolateral prefrontal cortex. The component could be a metacognitive process that monitors operations performed on working memory, or a process responsible for the maintenance of the goal structure needed to guide processing in any task (see, e.g., Cohen *et al.*, in press). Determining which of these, or other, alternatives is correct is high on the agenda of those who study the neural basis of working memory in hope of elucidating its role in higher-level cognition.

REFERENCES

Atkinson, R. C., & Shiffrin, K. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 2). New York: Academic Press.

Awh, E. (1996). *Attention-based rehearsal in spatial working memory.* Unpublished Ph.D. dissertation, University of Michigan, Ann Arbor, MI.

Awh, E., & Jonides, J. (in press). Spatial selective attention and spatial working memory. In R. Parasuraman (Ed.), *The attentive brain.* Cambridge, MA: MIT Press.

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

Awh, E., Smith, E. E., & Jonides, J. (1995). Human rehearsal processes and the frontal cortex: PET evidence. In J. Grafman (Ed.), *Structure and functions of the human prefrontal cortex.* New York: New York Academy of Sciences.

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

Baddeley, A. D. (1992). Working memory. *Science,* **225,** 556–559.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *Recent Advances in Learning and Motivation, Vol. VIII.* New York: Academic Press.

Braver, T. S., Cohen, J. D., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. *Neuroimage,* **5,** 49–62.

Broadbent, D. E. (1958). *Perception and communication.* London: Pergamon Press.

Carpenter, P. A., Just, M. A., & 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., Forman, S. D., Braver, T. S., Casey, B. J., Servan-Schriber, D., & Noll, D. C. (1994). Activation of prefrontal cortex in a non-spatial working-memory task with functional MRI. *Human Brain Mapping,* **1,** 293–304.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Noll, D. C., Jonides, J., & Smith, E. E. (in press). Temporal dynamics of brain activation during a working memory task. *Nature.*

Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (in press). A computational approach to prefrontal cortex, cognitive control and schizophrenia. Recent developments and current challenges. *Philosophical transactions of the Royal Society, Biological Sciences.*

Corbetta, M., Miezin, F., Shulman, G. L., & Petersen, S. E. (1993). A PET study of visualspatial attention. *Journal of Neuroscience,* **13** (3), 1202–1226.

Daneman, M., & Carpenter, P. A. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior,* **19,** 450–466.

Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience,* **4,** 2051–2062.

D'Esposito, M., Detre, J. A., Alsop, D. C., Shin, R. K., Atlas, S., & Grossman, M. (1995). The neural basis of the central executive system of working memory. *Nature,* **378,** 279–281.

Donders, F. E. (1868). On the speed of mental processes. Reprinted in *Acta Psychologica,* **30** (1969), 412–431.

Friston, K. J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. J. (1991). Comparing functional (PET) images: The assessment of significant change. *Journal of Cerebral Blood Flow and Metabolism,* **11,** 690–699.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology,* **61,** 331–349.

Fuster, J. M. (1995). *Memory in the cerebral cortex.* Cambridge, MA: MIT Press.

Fuster, J. M., & Jervey, J. P. (1981). Inferotemporal neurons distinquish and retain behaviorally relevant features of visual stimuli. *Science,* **212,** 952–955.

Goldman-Rakic, P. S. (1987). Circuitry of primate pre-frontal cortex and regulation of behavior by representational memory. *Handbook of physiology, Section 1,* **5,** 373–417.

Grasby, P. M., Frith, C. D., Friston, K. J., Bench, C., Frackowiak, R. S. J., & Dolan, R. J. (1993). Functional mapping of brain areas implicated in auditory-verbal memory function. *Brain,* **116,** 1–20.

Jonides, J. (1995). Working memory and thinking. In E. E. Smith & D. Osherson (Eds.), *Invitation to cognitive science: Thinking* (Vol. 3, 2nd ed. pp. 215–265). Cambridge, MA: MIT Press.

Jonides, J., Schumacher, E. H., Smith, E. E., Lauber, E. J., Awh, E., Minoshima, S., & Koeppe, R. A. (in press). Working-memory-load affects regional brain activation as measured by PET. *Journal of Cognitive Neuroscience.*

Jonides, J., & Smith, E. E. (in press). The architecture of working memory. In M.D. Rugg (Ed.), *Cognitive neuroscience.* London: University College Press.

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

Just, M. A., & Carpenter, P. A. (1992). A capactiy theory of comprehension: Individual difference in working memory. *Psychological Review,* **99,** 122–149.

Kolb, B., & Whishaw, I. Q. (1995). *Fundamentals of human neuropsychology* (4th ed.). New York: Freeman.

Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Maljkovic, V., Weise, S. B., Chabris, C. F., Hamilton, S. E., Rauch, S. L., & Buonanno, F. S. (1993). Visual mental imagery activates topographically organized visual cortex: PET investigations. *Journal of Cognitive Neuroscience,* **5,** 263–287.

Kosslyn, S. M., Thompson, W. L., Kim, I. J., Rauch, S. L., & Alpert, N. M. (1995). Individual differences in cerebral blood flow in area 17 predict the time to evaluate visualized letters. *Journal of Cognitive Neuroscience,* **8,** 78–82.

Kulpe, O. (1895). *Outlines of psychology.* New York: MacMillan.

McCarthy, R. A., & Warrington, E. K. (1990). *Cognitive neuropsychology: A clinical introduction.* San Diego: Academic Press.

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in Neuroscience,* **6,** 414–417.

Owen, A. M., Doyon, J., Petrides, M., Evans, A. C., & Gjedde, A. (1994). The neural mediation of high level planning examined using positron emission tomography (PET). *Society for Neuroscience Abstracts,* **20,** 353.

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

Petrides, M., Alivisatos, B., Evans, A. C., & Meyer, E. (1993a). Dissociation of human mid-dorsolateral from posterior dorsolateral frontal cortex in memory processing. *Proceedings of the National Academy of Science USA,* **90,** 873–877.

Petrides, M., Alivisatos, B., Meyer, E., & Evans, A. C. (1993b). Functional activation of the human frontal cortex during performance of verbal working memory tasks. *Proceedings of the National Academy of Science USA,* **90,** 878–882.

Posner, M. I. (1978). *Chronometric explanations of mind.* Hillsdale, NJ: Erlbaum.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience,* **13,** 25–42.

Posner, M. I., Peterson, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive functions in the human brain. *Science,* **240,** 1627–1631.

Raichle, M. E. (1996). Interview. *Journal of Cognitive Neuroscience,* **8,** 189–195.

Reuter-Lorenz, P., Jonides, J, Smith, E. E., Hartley, A. A., Cianciolo, A., Awh, E., Marshuetz, C., & Koeppe, R. A. (1996). The effects of age on the neural substrate for verbal working memory. *Society for Neuroscience Abstracts,* **22,** 183.

Roy, C. S., & Sherrington, C. S. (1890). On the regulation of the blood supply of the brain. *Journal of Physiology, London* **11,** 85–108.

Rubenstein, J. S., Meyer, D. E., & Evans, J. (1994). *Executive control of cognitive processes in task switching.* Paper presented at the Annual Meetings of the Psychonomics Society, St. Louis.

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

Shallice, T. (1988). *From neuropsychology to mental structure.* Cambridge: Cambridge University Press.

Smith, E. E., & Jonides, J. (1994). Working memory in humans: Neuropsychological evidence. In M. Gazzaniga (Ed.), *The cognitive neurosciences.* Cambridge, MA: MIT Press.

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

Smith, E. E., Jonides, J., Koeppe, R. A., Awh, E., Schumacher, E. H., & Minoshima, S. (1995). Spatial vs. object working memory: PET investigations. *Journal of Cognitive Neuroscience,* **7,** 337–356.

Smith, E. E., Patalano, A., & Jonides, J. (in press). Alternative mechanisms of categorization. *Cognition.*

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

Sternberg, S. (1969) The discovery of processing stages: Extensions of Donders' method. In W. G. Koster (Ed.), *Attention and performance* (Vol. 11), *Acta Psychologica,* **30,** 276–315.

Talairach, J., & Tournoux, P. (1988). *A co-planar stereotactic atlas of the human brain: An approach to medical cerebral imaging.* New York: Thieme Medical Publishers.

Waugh, N. C., & Norman, D. A. (1965). Primary memory. *Psychological Review,* **72,** 89–104.

Wilson, F. A. W., O Scalaidhe, S. P., & Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science,* **260,** 1955–1958.

Zatorre, R. J., Halpern, A. R., Perry, D. W., Meyer, E., & Evans, A. L. (1996). Hearing in the mind's ear: A PET investigation of musical imagery and perception. *Journal of Cognitive Neuroscience,* **8,** 21–46.

(Accepted February 12, 1997)