Modules of Working Memory

John Jonides
Ching-Yune C. Sylvester
Steven C. Lacey
Tor D. Wager
Thomas E. Nichols

University of Michigan

Edward Awh

University of Oregon

Address Correspondence to:

John Jonides
Department of Psychology
University of Michigan
525 E. University
Ann Arbor, MI 48109-1109

Email: jjonides@umich.edu
Fax: 734-994-7157

Summary

Working memory is best conceived as a set of modules responsible for the storage of information for a brief period of time and for the manipulation of this information in the service of ongoing tasks. To date, there has been considerable evidence from behavioral studies of normal and brain-injured individuals implicating separable storage and rehearsal processes as well as separable processes for verbal and spatial information. However, little evidence has accumulated about the architecture of executive processes. The addition of neuroimaging evidence concerning the executive processes as well as processes of storage and rehearsal enhances the picture of working memory provided by behavioral data.

Working memory is often defined as the memory system responsible for the storage of limited amounts of information for brief periods of time. With so narrow a definition, one may wonder what role working memory plays in our overall cognitive lives -- lives that are concerned with solving problems, with inductive and deductive reasoning, with language production and comprehension in the service of communication, and with intelligent behavior in general, whether in humans or other animals. By now there is growing evidence that working memory is indeed critical to higher cognitive life: We know this from studies of the strong relationship between performance on working memory tasks and performance on a large range of other cognitive tasks (e.g., Salthouse, 1991). We also know that when the brain structures that mediate working memory are compromised by illness or injury, not only does working memory itself suffer, but deficits also pervade the cognitive skills that it supports (see, e.g., Shallice and Vallar, 1990). In short, understanding the mechanisms of working memory will have benefits not only for understanding the architecture of this isolated memory system, but also for understanding changes in a large repertoire of cognitive skills.

At present, we know a good deal about the architecture of working memory. Our knowledge derives principally from three sorts of empirical programs: behavioral studies of normal adults, behavioral studies of brain-injured patients, and neuroimaging studies of normal adults. Together, these sources of information are leading to the development of a comprehensive view of both the psychology of working memory and the underlying neural architecture that supports the psychology. What has become increasingly clear from the accumulated research is that working memory is best conceptualized not as a monolithic construct, but rather as a set of modules. One conceptualization is that the modules can be grouped along two dimensions. One dimension has to do with function--whether a module is involved in storage of information, in rehearsal of that information, or in manipulation of that information for some cognitive purpose. The other dimension concerns the nature of the information that is stored in working memory--whether verbal, spatial, or some other code.

To appreciate the modularity of the organization of working memory, let us begin with the theoretical framework first introduced by Alan Baddeley and his colleagues (1986, 1992). While not uncontroversial, the current version of this framework proposes that the storage of information in working memory is accomplished by a set of storage buffers, each responsible for a different sort of information – that is, the buffers are defined by the type of information they store. Each buffer has a rehearsal function associated with it to refresh the information stored there so that it can survive the normally short durations of unrehearsed memory traces. The contents of each of the buffers are then available to a set of executive processes that can manipulate the memorial representations in the service of some ongoing task, such as mental arithmetic, comprehending spatial directions, or reasoning.

To see how such a system might work, consider the processes required to solve a mental arithmetic problem such as:

74 x 12 = ?

First, of course, the problem itself must be stored in working memory until a solution is reached. In one way of solving the problem, the solver must attend to the “tens” digit of the “74” (i.e., “7”)and retrieve a rule or table from memory in order to multiply this by “12” to yield “840.” This intermediate solution must then be stored temporarily while attention is turned from it to the units digit of the “74.” Again, a multiplication rule or table must be retrieved from long-term memory so that the “4” can be multiplied by “12” to yield “48,” another intermediate solution that must be stored. Then the first intermediate solution, “840”, presumably being rehearsed in the background, must be retrieved and addition rules or tables also retrieved so that “840” can be added to “48” to yield the final answer of “888”. Of course, all this storage, retrieval and computation must be completed in the face irrelevant information in the environment that might interfere with performance. Even this simple arm-chair analysis of the processes involved in mental arithmetic reveals that the working memory mechanisms that are recruited to the task are both storage processes, and executive processes that coordinate operations performed on the stored information.

Of course, intuition suggests that while the processes involved in this sort of problem-solving involve arithmetic information, they are also heavily language-based. However, working memory extends to other information domains besides language. Consider another example to appreciate this. Suppose someone gives you directions from your home to the local grocery. She might tell you to make a left at your driveway, go to the second traffic light, make a right until you reach the gasoline station, make a left there to the elementary school on the right, proceed one block past the elementary school to the stop sign, make a left there and go 5 blocks to the grocery. Many people find that an effective strategy for storing such directions is to store a mental route that is described by the directions. That is, the listener would construct a spatial representation from the verbal information and use that to guide himself. To do so, one would have to encode the information in terms of spatial features (such as visualizing the directions left and right or creating images of the landmarks that are named), organize these spatial features in appropriate order, store the whole spatial representation of the route, and retrieve parts of it appropriately. Once again, this analysis suggests that working memory and executive functions play roles in problem-solving by mediating the storage of information and the manipulation of that information. Of course, this task may also require the use of long-term memory to retrieve familiar landmarks, information about directions, and so forth. Indeed, it may be that most of the route is stored in long-term memory, with only a portion of it activated as it is needed. In spite of these considerations, the task still clearly places heavy demands on working memory.

These examples nicely illustrate the intuition that working memory is characterized by distinguishing different kinds of processes (storage, rehearsal, and executive functions) and by different kinds of information (verbal and spatial in the examples given, but others as well, such as visual information that is not spatial). However, a proper theory of working memory must be built on more than intuitions, and there is by now a wealth of evidence supporting the architecture that is suggested by these examples. Rather than being comprehensive in reviewing this literature, we shall concentrate on experimentation from our own laboratory using both behavioral and neuroimaging techniques to offer evidence that is relevant. First, we shall review studies indicating that storage mechanisms for different kinds of information in working memory are separable from one another. Second, we shall show that storage can be separated from rehearsal for verbal information, and possibly for spatial information as well. Finally, we shall offer evidence that executive processes are best characterized not as a single controller, but rather as largely separable mechanisms which share some common neural underpinnings.

Verbal versus Spatial Storage

To show that working memory for different kinds of information recruits different brain mechanisms, one would like an experimental setting in which the same memory task can be performed on different types of information, with little involvement of executive processes; that is, the processes that should be the focus of the task are those involved in storage and rehearsal. One task which fits this requirement is the item-recognition task. In the item-recognition task, subjects are given a set of items to store for several seconds, after which a probe item is given, and subjects must indicate whether this item was a member of the memorized set. Notice that this task places little requirement on executive processes because there is no manipulation required of the stored information; instead, the task emphasizes the storage of the items, the rehearsal of those items, and retrieval processes necessary to decide if the probe had been presented as part of the memorized set.

The item-recognition task is nicely suited to studying working memory for different types of information because one can easily prescribe what items must be stored. In a pair of experiments, we have done just this, as illustrated in the left panel of 1 (Smith et al., 1996). In the Memory condition of a Spatial working memory task, subjects saw three dots at unpredictable locations on a screen that they had to store in memory for three seconds. Following this retention interval, a single location was probed, and subjects had to indicate whether this location was one of the three they had stored in memory. The comparable Verbal Memory task is indicated in the right panel of the figure. Here, subjects were presented four letters that they also had to store for three seconds, following which a single letter was presented, and they had to indicate whether this was one of the stored letters. Different groups of subjects participated in these two tasks while being scanned using positron emission tomography (PET). Appropriate control conditions for each memory task are also shown in Figure 1. The control tasks were designed so that subjects were presented with similar perceptual displays and had to make similar matching judgments, but the memory requirement in each control condition was minimal. Consequently, contrasting the activations of the Spatial and Verbal Memory conditions with each of their control conditions should yield activations due to the storage and rehearsal of spatial and verbal information respectively, and not due to encoding operations or response processes.

Behaviorally, subjects were quite accurate in these tasks. Of interest is that their response times for the Memory conditions exceeded the response times for the Control conditions. This is consistent with the assumption that the Memory conditions required processes in addition to those engaged in the Control conditions, presumably processes of storage and rehearsal.

The brain activations reveal a pattern that indicates a dissociation between verbal and spatial working memory. These activations are shown in Figure 2, with the verbal condition shown on the top row and the spatial condition on the bottom. Perhaps the most obvious feature of the activations shown in the figure is that they differ by hemisphere, with greater right-hemisphere activation in the spatial condition and greater left-hemisphere activation in the verbal condition. This difference is of great interest in describing the architecture of working memory because, by virtue of the design of the experiment, it represents largely the storage and rehearsal components of the task, not those due to encoding or retrieval. Thus, the difference in hemispheric asymmetry in these activations indicates that there is a difference in the mechanisms responsible for maintenance in working memory based on the type of information being maintained.

Beyond this gross difference, there are also more detailed features of the activations that merit comment. In the verbal task, the major sites of activation are in inferior frontal gyrus near Broca's area, in premotor cortex in the supramarginal gyrus of posterior parietal cortex, and in superior parietal lobule, all concentrated in the left hemisphere. As we shall see below, the activation in Broca's area can be distinguished from the activation in parietal cortex in such a way as to associate the former with rehearsal and the latter with storage. For the spatial task, the most noticeable activations appear in prefrontal cortex in the region of superior frontal sulcus and inferior frontal gyrus, as well as in extrastriate cortex in the occipital lobe. The functions of these regions have not yet been clearly described, compared to those for the verbal task, but there is some evidence, reviewed below, that the extrastriate activation reflects the operation of a spatial rehearsal process, possibly involving the use of covert spatial orienting. The dissociation revealed by these data has been replicated by others, indicating the robustness of the finding that the neural circuitries for storage in spatial and verbal working memory are different from one another (see, e.g., Paulesu et al., 1993; Courtney et al., 1996). Beyond this, there is also some evidence that information about object form recruits yet another set of storage mechanisms (Courtney et al, 1996). It should also be noted that the different circuitries that appear to be involved in working memory are not simply defined by input modality; the distinction seen in this task between visually presented spatial information versus visually presented non-spatial information makes this point. In addition, there is evidence from another PET study that verbal information entered into working memory by ear or by eye makes little difference to the storage mechanism that is used (Schumacher et al., 1996). Thus, what appears to be the defining characteristic of the different storage mechanisms is the information that is stored, not the way that information first enters the system.

Rehearsal

The information stored in working memory is fragile; without being refreshed, it will decline in strength in a matter of seconds and be unavailable for retrieval. To keep the information active, one must engage in some sort of rehearsal process that recycles and refreshes the traces, thereby mitigating the effects of decay and interference. This view of working memory highlights the distinction between processes responsible for storage of information and processes responsible for rehearsal, and there is behavioral evidence to support this distinction both from normal and brain-injured subjects (e.g., Basso, Spinnler, Vallar, & Zanobio, 1982; Longoni et al., 1993). The behavioral evidence is not itself completely satisfying, however, largely because it is difficult to investigate the effects of rehearsal separately from the effects of storage, and because patients with damage to one system often show carryover effects in other systems. We have documented the separability of storage and rehearsal using positron emission tomography (PET) measurements of a task that has been used widely to study working memory, the n-back task (Awh et al., 1996).

In this task, subjects are presented a sequence of single-letters at a fairly leisurely pace (once every 3 sec in the experiment to be described). In the 2-back version of the task, when each letter is presented, a subject must decide whether it matches in identity the letter presented two items back in the sequence. So, for example, if the letters P, N, P, R, J, R, L, D, D were presented, subjects should respond positively to the third and sixth letters in the sequence ("P" and "R"). Note that subjects should not respond to the last letter ("D"). This is because while the third letter matches the first, and the sixth matches the fourth, the last letter matches the one just before it, not the one two items before it. This is a demanding task that requires focused attention, but with a bit of practice, subjects become quite accurate, scoring above 90% in accuracy. The task clearly requires storage of each letter in memory at least until two more letters are presented, and to maintain this storage, rehearsal is also required otherwise the interfering effects of successive letters would cause the memory traces to degrade substantially. In addition, of course, unlike the item-recognition task, the 2-back task requires considerable executive processes that are responsible for updating the contents of working memory and for inhibiting positive responses to letters that do not match the letter two items back in the series.

In this experiment, our intention was to focus on the processes required for storage and for rehearsal, not the executive processes about which we shall comment below. In order to isolate storage and rehearsal, we also had subjects participate in two control conditions. In a "Search" control, a single target-letter was presented to subjects at the beginning of a series of single-letter presentations (as in the 2-back condition), and this letter served as the target for the entire series. When subjects saw it, they were to respond positively; otherwise, they were to respond negatively. Compared to the 2-back task, the Search control has a smaller storage and rehearsal requirement, so contrasting activations in the Search condition to those in the 2-back condition should reveal regions responsible for storage and rehearsal. In the "Rehearsal" control condition, subjects were again presented a series of single letters at the same pace as in the 2-back task. When each letter was presented, they were instructed to rehearse it silently until the next one appeared, at which time they were to rehearse that one, and so on. This condition has an even smaller storage requirement than the Search condition, but places substantial demand on rehearsal. Thus, subtracting the Rehearsal control from the 2-back task should substantially subtract out rehearsal processes, but leave intact activations due to storage.

Figure 3 displays the activations from the PET measurements that resulted from subtracting the Search and Rehearsal controls respectively from the 2-back task. The data are displayed for 4 comparable horizontal brain slices from each subtraction,revealing the relative activations. Consider first the activations shown in the left panel for the subtraction of the Search control from the 2-back task. This panel reveals prominent bilateral activations in parietal, premotor and supplementary motor cortex, inferior frontal gyrus on the left, as well as cerebellum. Contrast this with the subtraction of the Rehearsal control from the 2-back task, shown in the right panel. The major difference between this subtraction and that with the Search control is that the activation in inferior frontal gyrus is dramatically reduced. By the logic of the subtraction method, this difference leads to the conclusion that the left inferior frontal gyrus is a major contributor to verbal rehearsal processes. This makes sense when one considers that this same region is heavily involved in overt articulation as well, suggesting that the same region that is critical to speech production is co-opted for internal rehearsal. Note also that the activation in parietal cortex in the subtraction of the Rehearsal control from the 2-back condition is somewhat larger than in the subtraction of the Search control from the 2-back condition. This is as it should be if these parietal activations reflect storage processes because the storage requirements of the Search control are larger that those of the Rehearsal control.

The concept of rehearsal for verbal material has much intuitive appeal to it because many subjects engaged in working memory tasks do have the introspection that they devote effort to “talking to themselves” to keep traces fresh. If the storage/rehearsal architecture is a general one, though, it ought to apply to other types of material stored in working memory as well. Consider again storage of spatial material, as required in the item-recognition task described above. If this material is not rehearsed during a retention interval, the traces of the locations that were encoded will quickly fade from memory. Thus, one might propose that there is rehearsal for spatial information similar to that for verbal information; more generally, one might propose that working memory for any sort of information requires a rehearsal process. What is interesting about this proposal is that the nature of rehearsal must differ for different types of information. For example, it is not effective to use a verbal rehearsal strategy for the dot-locations of the spatial item-recognition task because the probe may appear in a location that is categorically similar to one of the target locations (e.g., upper left or lower right third of the screen), thereby activating the same verbal code, but it may nevertheless not match a target location. What is required for each type of information is a rehearsal strategy that is tailored to the type of information in question.

What might this strategy be for spatial information? One interesting possibility is raised by the striking similarity in brain regions involved in spatial working memory and the allocation of attention to places in space, as reviewed by Awh and Jonides (1998; 2001). This analysis suggests that in both humans and other animals, there is substantial commonality in the brain mechanisms involved in the two tasks. This commonality raises the hypothesis that rehearsal of spatial information may involve an allocation of spatial attention to the specific locations that are being stored. We have tested an implication of this: If spatial rehearsal engages an internal attention-allocation mechanism for a rehearsed location, then there should be evidence of improved processing at that location if a visual discrimination task is inserted during the retention interval of a spatial memory task (Awh et al., 1998). The task is illustrated in the top panel of Figure 4. A letter was presented briefly in a location followed by a retention interval of 5 sec. Following this interval, another letter appeared at a location on the screen, and subjects made a judgment about whether this letter matched the earlier presented target. Half the subjects judged whether the letters matched in their spatial location, regardless of the identity of the letters, and half judged whether the letters matched in identity regardless of spatial location.

What is critical in the experiment is what transpired during the retention interval. During this interval, subjects made speeded discrimination judgments about the left or right orientation of a nonsense symbol that was presented on the screen, as shown in Figure 4a. This figure sometimes appeared in the location that subjects were storing in memory and sometimes in another location. The critical prediction was that when subjects were storing spatial information, if the figure appeared at the memorized location, it would be discriminated faster than if it appeared at a non-memorized location. However, if the memory task involved letter identity and not position, there would be no difference in discrimination speed regardless of where the figure appeared during the retention interval. As Figure 4b shows, just this pattern of response times emerged. What these data reveal is that some processes engaged by the memory task have a carryover effect on a visual discrimination task (during the retention interval) that should be influenced by the allocation of visual attention to a spatial location. Thus, these data suggest that spatial rehearsal engages the same attention mechanism used for visual discrimination. What the data do not indicate, however, is whether this spatial attention mechanism is functional for the memory task itself.

To address this issue, we conducted another experiment that used the same main memory task, working memory for a single spatial location, as shown in Figure 5 (Panel a) (Awh and Jonides, 2001). Again, a visual discrimination task intervened during the retention interval (dual-task), but in this case the discrimination task involved judgments about hues (whether a color patch was red or blue) rather than letter or location. In the condition schematized at the top of the panel, the color discrimination had to be made to a small patch presented eccentrically on the screen, and so attention had to be directed to it in order to make that judgment. In the task shown just below this, the color patch was large and occluded all of the potential memory locations. Thus, in this condition subjects could discriminate the color of the large patch without shifting attention away from the memorized location. There was a significantly larger decline in memory accuracy for the target spatial location when subjects had to discriminate the color of the small patch during the retention interval (i.e., when the color discrimination required a shift of attention away from the memorized location) than when subjects discriminated the color of the large patch (as shown by the right two points in Figure 5b). To be sure that this effect was a function of the intervening shift of attention to the target color, a control condition was included in the experiment in which the colors were presented during the retention interval, but no judgment had to be made about them; subjects just had make judgments about the target spatial location in this control condition. The data from this control (single-task) condition are shown in the left two points of Figure 5b; these data reveal that without the dual-task requirement, memory performance was approximately equivalent when the large and small color patches were presented during the retention interval. Note also that the experiment yielded a replication of the effect of attentional allocation on the intervening color discrimination task during the dual-task, as shown in Panel 5c. When the color patch was large and presented in the center of the screen, having this judgment made in the context of a memory task did not make performance worse compared to when it was made alone. However, when the spatial memory task was required, then performance suffered for the color judgment on the small color patch, presumably because attention was not allocated to the location of the patch.

The data from these two experiments taken together suggest that spatial working memory and spatial attention share a mechanism in common, a mechanism that for spatial working memory operates during the retention interval of a memory task. Recall from the item recognition experiments that one of the regions activated by spatial storage in working memory was extrastriate cortex. It is now well-documented that activity in visual cortex can be modulated by spatial selective attention, with larger visually-evoked responses in the brain regions that process the attended locations. Following the possible analogy between spatial attention and spatial working memory, perhaps rehearsal in spatial working memory is related to modulation of activity in extrastriate cortex, not necessarily in the service of better perception of spatial information, but in the service of better retention of that information. To address this issue, we conducted an experiment in which functional magnetic resonance imaging (fMRI) was used to measure whether extrastriate cortex was modulated in its activity when subjects were engaged in a spatial working memory task compared to a verbal working memory task (Awh et al., 1999). In particular, we sought to test whether rehearsal of a spatial location produced increased activation of extrastriate cortex contralateral to that location, just as attention to a visual stimulus produces activation in occipital cortex contralateral to the position of that stimulus.

The paradigm is illustrated in Panel a of Figure 6. Three target memoranda were presented sequentially and followed by a retention interval of 7 sec during which time a checkerboard was flashed to both visual fields. The memoranda were false-font characters in either the left or right visual fields. After the retention interval, a single false-font character was presented, and subjects had to judge whether it appeared in a location marked by one of the target characters. We measured the activation in extrastriate cortex due to the flashing checkerboard to see whether these activations were larger in the field contralateral to the memorized locations. To control for the fact that the stimuli themselves were presented unilaterally, in a control task, the same subjects engaged in a verbal (letter) memory task in which memoranda were again presented unilaterally, but the memoranda were letters rather than false-font characters. Subjects then had to judge whether a probe letter presented after the retention interval matched in identity one of the target letters, regardless of position. For individual subjects, the regions of occipital cortex that were activated by the checkerboard were mapped using data from a separate control task so that we could examine individually whether these areas were modulated by the spatial rehearsal task compared to the verbal rehearsal task. Panel 6b shows that the spatial rehearsal task yielded greater activation in these contralateral voxels than the verbal task for all the slices measured in posterior cortex. Thus, it appears that spatial rehearsal in working memory leads to a modulation of activity in extrastriate cortex in just the same way that spatial selective attention does. This supports the view that there is significant functional overlap between these two processes.

To sum up our discussion of storage and rehearsal in working memory, we have found evidence illustrated by research in our laboratory, but confirmed by work in other laboratories, of two central features. One is that working memory storage appears to be mediated by different mechanisms as a function of the type of information stored. The second is that rehearsal provides support for continued storage of information beyond the brief period during which it would be viable without intervention. Furthermore, the particular mechanisms of rehearsal are tied to the type of stored material. For verbal material, mechanisms ordinarily responsible for the production of speech play a central role in internal recycling of information. For spatial material, mechanisms responsible ordinarily for visual selection are co-opted to the task of maintaining internal spatial codes. These two examples illustrate nicely how the nervous system manages to harness single mechanisms for multiple tasks based on the computational competence of the mechanisms. We speculate more generally that there may be many memory processes that ride piggyback on posterior mechanisms that evolved for the purpose of processing sensory information from the outside world (see, e.g., Miller et al., 1996, for a similar point about working memory for information about objects).

Executive Mechanisms

We turn now to the final set of modules important for working memory, the executive functions responsible for manipulating information. There has been a good deal of theoretical debate about the nature of executive processes that has focused on one issue: Are executive processes of a piece, or are there multiple such processes each largely different from the others but acting in concert to accomplish various task goals?

One might think that the singular view, on the face of it, must be a straw man: Any task that makes heavy use of executive processes, such as mental multiplication or the n-back task discussed above, appears to recruit many different processes. Recall the processes of mental arithmetic and the n-back task to illustrate this point. The mental arithmetic task requires at least two seemingly quite different kinds of attention-switching processes in the service of manipulating information in working memory. One of them is responsible for switching between different internal representations of parts of a problem (e.g., the "units" digit versus the "tens" digit) in order to perform some operation on it. By contrast, at some point in the problem, one must also switch between one mental computation (multiplication) and another (addition). While both of these require shifts of attention, the processes that underlie each of the shifts seem, on the face of it, to be different from one another.

Now consider the n-back task, which seems to feature a set of seemingly very different kind of manipulation or executive processes, those involved in inhibitory functions. One such kind of inhibition in the 2-back task is the ability to inhibit a positive response to an item that appeared just previously in a sequence (i.e., 1-back) because that item, although familiar, is not a 2-back match. Note also that this task requires another kind of inhibition – the inhibition of the item that is two items back in the sequence so that a new letter can be entered into memory to replace the current letter. For example, in the sequence, "L, P, L, R", when the second L appears and is matched to the item two letters back, the first L must then be discarded from memory in order for the new letter, R, to enter and be marked as the current letter to be matched against its two-back candidate, P. Thus, the n-back task relies more on processes responsible for inhibiting irrelevant representations rather than those responsible for switching between internal representations or internal computations. In the face of such seemingly different sorts of switching and inhibitory processes, how could one argue that executive processing is "of a piece?"

In fact, such an argument is not a straw man if one conceptualizes all executive processes as variants of attentional allocation mechanisms. In the case of switching between representations or operations in mental arithmetic, perhaps the central mechanism that is required is one that activates the alternative representation or operation, and suppresses the currently active one. In the case of inhibition in the n-back task, an analysis in terms of attentional allocation may also be appealing. For example, perhaps what is involved in avoiding an incorrect positive response to a match with an item that is just one letter back in the series is allocation of attention to a representation tied to the episodic tag of when an item appeared in a series, and inhibition of the potency of a familiarity representation. Similarly, to rid memory of an item in a series to make way for a new item, perhaps what is involved is an activation of the representation of the newest representation, and suppression of the representation of the oldest item of the set. This view casts executive processes entirely in terms of attentional allocation, and as such, it has a respectful place among theories of executive processing. Indeed, papers by Baddeley and his colleagues (1986, 1992), Norman and Shallice (1986), and others are all examples that make a good theoretical case for a singular view of what has been called the "central executive," a term that suggests a singular vision of this sort of processing.

By contrast, consider an alternative view of executive processes, in which individual processes are different computationally from one another, and therefore could be mediated by quite different brain mechanisms. Using our examples of mental arithmetic and the n-back task, we might suppose that there are several quite different processes involved. In the case of mental arithmetic in which attention must be switched from the representation of one digit to the representation of another, this switching might be accomplished by a mechanism that changes the activation values of the two memory representations, bringing one to the forefront and relegating the other to the background. This sort of change in activation levels may be the sort of mechanism needed when two representations have equivalent status and changes need to be made rapidly between them to accomplish some task, such as multiplication of numbers. To change from one mental computation to another (say, multiplication to addition), what may be required is a switch in what rules or tables are retrieved from long-term memory. Thus, even though these two kinds of processes may be cast similarly by using the term "switching," what is actually involved may be computationally quite different.

The n-back task may involve yet different executive processes from those used in mental arithmetic. Perhaps removing the oldest item from working memory and adding a new item does involve some sort of attentional allocation scheme that changes the relative activations among items in memory. By contrast, preventing a positive response to an item that was just presented and that matches the current item (a 1-back false alarm) might involve an inhibitory process that blocks the prepotency of highly familiar items. This way of looking at the difference between switching and inhibition suggests that the two should be dissociable behaviorally. Further, if they are dissociable, then there should be evidence of different brain mechanisms responsible for their mediation.

We have recently conducted a combination of behavioral and neuroimaging experiments to determine whether switching of attention between representations in working memory and inhibiting attention to a prepotent response are separable processes (Sylvester et al., submitted). Our research began with a behavioral task illustrated in Figure 7. Subjects were presented with a sequence of arrowheads that pointed left or right, each presented until the subject made a response. One of the tasks was to keep track of the number of left-facing arrows and the number of right-facing arrows in a sequence of 8 to 11 arrows. After each sequence of arrows, subjects were probed about their counts to assess accuracy. Note, as shown in the figure, that the sequence of arrows yielded two types of trials of interest: those on which a succession of two arrows pointed in the same direction so that subjects did not have to switch counters versus those on which a succession of two arrows pointed in different directions, necessitating a switch in counters in order to update the counts. This task is modeled after one introduced by Garavan (1998) who showed that the response time to each stimulus that had to be counted depended on whether it indicated the same counter as the previous stimulus (non-switch trials) or whether it indicated a different counter from the previous stimulus (switch trials). Switches took longer. The other task that faced subjects was to respond with a manual keypress to the presentation of each arrow. In one condition, subjects' responses were compatible with the direction in which the arrow pointed (i.e., a left keypress to the presentation of a left-facing arrow and a right keypress to the presentation of a right-facing arrow). In the other condition, the assignment of responses to arrows was reversed so that the responses were incompatible with the directions of the arrows. A large corpus of previous research using manipulations of stimulus-response compatibility has shown that there is a reaction time cost when the required response is incompatible with the prepotent response to a stimulus. Compatible and incompatible blocks were alternated across each 8-11 arrow sequence. The main issue underlying this experiment was whether the cost in switching between counters and the cost in responding incompatibly are dependent on one another as they might be if there is a single mechanism underlying them; alternatively, these two performance costs might be statistically independent of one another if there are two separate executive mechanisms involved in mediating the two operations.

The behavioral results from the experiment are quite clear. Very early in performance, as can be seen in the left side of Figure 7b, one can see the interaction of switching and compatibility reaction time costs. This interaction was initially over-additive in nature (accuracy in both tasks is near ceiling, so response time is the dependent measure of choice to examine performance on this task). Note, though, that the small over-additivity was ephemeral, converging to zero after several 16 block runs of trials, after which, as shown on the right side of Figure 7b, the two factors affected performance independently. These results lead to the following hypothesis: That there are two separable mechanisms, one mediating switches of attention to internal representations and another mediating rules that map stimuli onto incompatible responses (see Sternberg, 1969, for the rationale that underlies this hypothesis from response times). While these two mechanisms are separate, early in practice their independent operation is overcome by a common resource limit that is placed on performance by the unfamiliarity of the task, so it appears as if they interact; but this effect soon dissipates.

The behavioral results are not unambiguous, however. It may be that there is a single mechanism that mediates counter-switching and stimulus-response rule assignment, and that this mechanism (which speeds up with practice) operates on the two tasks in turn. If this were so, it could result in a pattern of seeming independence between the two performance costs, as was obtained. To gather further evidence of the independence of the two mechanisms, a functional MRI experiment was conducted to examine what brain areas are engaged by the two experimental factors. In this experiment, the counter-switching and stimulus-response compatibility tasks were wholly separate from one another. For both tasks, the stimuli consisted of a sequence of arrows presented one at a time, as in the behavioral experiment. The counter-switching manipulation was tested by having subjects keep track of two counters (left and right arrows) in two types of blocks. In one type of block, subjects had to switch between counters relatively frequently (70% of trials); in the other type, they had to switch relatively infrequently (20% of trials). On each trial in these blocks, subjects responded to each arrow with a double keypress using both left and right index fingers, so the responses to each arrow were identical, with no compatibility variation. The compatibility manipulation was tested in yet different blocks, contrasting a block of trials in which subjects responded compatibly to each arrow versus another block of trials in which they responded incompatibly, neither of which blocks involved any counting. Thus, this experiment examined the effects of each variable separately from the effects of the other, but within the same subjects.

One potential concern with the imaging data was that both experimental tasks might involve more overt or intended eye movements than their controls, and that mechanisms controlling these eye movements might produce activations that were not of interest. To rule out this possibility, a saccade-control task was included in the experiment. For this control task, we instructed subjects to execute saccades to a series of stimuli presented on a screen. Saccade-related activity was contrasted with a control condition in which subjects maintained fixation, resulting in bilateral superior frontal and superior parietal activations. In order to examine the areas of activation uncontaminated by the activations due to saccades, we subtracted out the saccade-related activations from those due to switching and compatibility in further analyses.

Having subtracted out the effects of eye movements, we proceeded to categorize the areas of activation observed in the switching and compatibility contrasts as being common to both switching and compatibility, or uniquely associated with one of these two variables. We created regions of interest consisting of all the activations due to both switching and compatibility, and determined which voxels within this ROI were active in both tasks (see Figure 8). These included bilateral superior parietal cortex, superior colliculus, anterior cingulate, left middle frontal gyrus and bilateral premotor cortex. The bilateral parietal cortex and superior colliculus have been implicated in selective attention (Buchel et al., 1998; Casey et al., 2000), the dorsolateral prefrontal cortex is thought to be involved in maintaining contents of working memory that may need to be manipulated, while the anterior cingulate may involved in detecting or responding to the conflict that arises from a competing dominant response (Gehring & Knight, 2000; MacDonald, Cohen, Stenger and Carter, 2000; Jonides et al., in press). The bilateral premotor cortex may be involved in the inhibition of a prepotent response in the response compatibility task. In the switching task, it is possible that subjects are actively delaying responses on switch-trials relative to non-switch trials since on switch-trials they must complete the counter-switch and counter-updating processes before they can make their motor responses, which may necessitate inhibition of a response, indexed by premotor activation. Overall, then these activations reveal that there are certain processes, such as detection of conflict and inhibition of a prepotent response that seem to be common to these two tasks and that reflect the recruitment of common executive mechanisms.

In contrast to these regions of common activation, there were also areas of activation that were better described as being unique to each of the tasks. We discovered regions of activation distinctive to the switching and compatibility variables by comparing switch-related activations to compatibility-related activations in the same regions of interest we used to assess common activations across the tasks. Using paired t-tests, we found that the counter-switching task yielded significantly greater activation in bilateral extrastriate cortex, left posterior superior parietal cortex, superior colliculus, left dorsolateral prefrontal cortex, and anterior cingulate (see Figure 9a). Together with the analysis of common activations, this suggests that there is unique activation in the extrastriate and posterior superior parietal regions for the counter-switching task, and greater activation due to switching than response-inhibition in the other areas (even those these areas are active at a lower level in the stimulus-response compatibility task). The posterior superior parietal cortex appears to be involved in attention-switching (Dove, Pollmann, Schubert, Wiggins & von Cramon, 2000), while the extrastriate activation may be a result of the use of mental imagery to represent the counters (see, e.g., Kosslyn et al., 1993, for evidence of the involvement of occipital cortex in mediating visual imagery). Supporting this hypothesis are the reports of several subjects who stated that their representations of the counters had a spatial quality to them (consistent with the left-right difference in the stimuli that mark the counters). These distinctive activations are an indication that the switching task recruits processes that are task-specific (such as the occipital activation that may reflect the recruitment of imagery processes) and additional executive processes (such as may be reflected in the dorsolateral prefrontal activation) that are not recruited at all or as much by the compatibility task.

Areas of activation that were better described as uniquely associated with response incompatibility were the bilateral superior parietal and supplementary motor area, as well as the right frontopolar cortex (see Figure 9b). It is important to remember that the parietal and superior frontal cortex activation observed in this analysis are not a result of saccade-related activity in that we subtracted this out prior to comparing activation due to switching and compatibility. Rather, these areas may be involved with motor response inhibition and response selection (Rubia et al, 2000). The frontal area (BA 10) may be involved in the maintenance and monitoring of a subgoal (i.e. "respond opposite") before the correct response can be made (Braver & Bongiolotti, 2002). So, the activations that are uniquely emblematic of the compatibility task are those that may reflect motor-related processes, ones that would not be needed in the switching task.

What do these commonalities and differences between tasks tell us about executive processes? One suggestion is that there is, indeed, a common executive mechanism involved in the allocation of attention in both the counter-switching and response inhibition task. This common mechanism, which is parietal in location, controls the allocation of attention—whether to an internal representation or to a mapping-rule. However, attentional allocation alone cannot account for the execution of both task-switching and response-inhibition processes. There are also separable mechanisms that mediate the switching of attention and the inhibition of a prepotent motor response. For the counters, this involves the actual switching of attention from one counter to another, and may be controlled by a region in superior parietal and dorsolateral prefrontal cortex. For the stimulus-response mapping, it involves the maintenance of a general task goal, controlled by frontopolar cortex, as well as motor programming operations regulated by supplementary motor and parietal areas that allow one to inhibit a prepotent motor response and select an alternative response.

Conclusions

When looked at in this way, the issue of whether executive mechanisms are unitary in nature or are composed of multiple types dissolves into the particulars of the mechanisms needed for any task. As models such as EPIC (Meyer and Kieras, 1997) suggest, what may prove to be the most productive exercise in understanding executive processes is a detailed modeling of what mechanisms are needed to yield behavioral performance. What our neuroimaging data suggest is that some of these mechanisms may be common to tasks that require executive control and others may be quite different. Further progress in understanding these mechanisms will come from mapping them in more and more tasks to understand when there are similarities among mechanisms and when there are differences.

We have come to understand a good deal about working memory by examining a combination of behavioral and neuroimaging data. What has emerged from this examination is a picture of a complex system with a modular organization. This organization appears, as originally proposed by Baddeley (1986, 1992), to be well-characterized by differentiating between processes responsible for maintenance and those responsible for information-manipulation. Maintenance appears to be mediated by sites in which information is stored and rehearsed. Storage and rehearsal processes, furthermore, are themselves not singular in form; they vary with the kind of code being maintained, and this variation is reflected in the brain mechanisms that are activated by different kinds of information as well as by the effects of brain lesions on particular deficits in working memory (see Jonides et al., 1996, for a discussion). Executive processes have multiple mechanisms underlying them as well. What appears to be a common thread among these mechanisms is the need to allocate or withdraw attention successively between alternative representations. How this is accomplished, though, depends on the particulars of the task in question. It may also depend on the type of information being processed (e.g., verbal versus spatial), although at this time, there is insufficient data to address whether this is so. With this sort of componential view of working memory developing, what is now needed is more extensive research on the particulars of each component, working toward a comprehensive model of the multiple modules of working memory.

Acknowledgement

Preparation of this manuscript was supported by a grant from the National Institute of Mental Health to the University of Michigan (John Jonides, Principal Investigator).

References

Awh, E., Jonides, J., Smith, E.E., Buxton, R.B., Frank, L.R., Love, T., Wong, E.C., & Gmeindl, L. (1999). Rehearsal in spatial working memory: Evidence from neuroimaging. *Psychological Science, 10*(5), 443-437..

Awh, E., Jonides, J., & Reuter-Lorenz, P.A. (1998). Rehearsal in Spatial Working Memory. *Journal of Experimental Psychology: Human Perception and Performance. 24(3)*, 780-790.

Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and working memory. *Trends in Cognitive Sciences,5*(3), 119-126.

Awh, E., & Jonides, J. (1998) Spatial Selective Attention and Spatial Working Memory. *The Attentive Brain.* Parasuraman, R. (Ed.), pp. 353-380, Cambridge, Mass: M.I.T. 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*(1), 25-31.

Baddeley, A. (1986). *Working Memory.* Oxford, England: Clarendon Press.

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

Basso, A., Spinnler, H., Vallar, G., & Zanobio, M.E. (1982). Left hemisphere damage and selective impairment of auditory verbal short-term memory: A case study. *Neuropsychlogia, 20*, 263-274.

Braver, T. S., & Bongiolatti, S. R. (2002). The role of frontopolar cortex in subgoal processing during working memory. *Neuroimage*, published electronically 1/22/02.

Buchel C., Josephs O., Rees G., Turner R., Frith C. D. & Friston KJ. (1998). The functional anatomy of attention to visual motion. A functional MRI study. Brain, 121, 1281-94

Casey, B.J., Thomas, K. M., Welsh, T. F., Badgaiyan, R. D., Eccard, C. H., Jennings, J. R., and Crone, E. A. (2000). Dissociation of response conflict, attentional selection, and expectancy with functional magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 97, 8728-8733.

Courtney, S.M., Ungerleider, L.G., Keil, K., & Haxby, J.V. (1996). Object, and spatial, visual-working memory activate separate neural systems in human cortex. *Cerebral Cortex,6,* 39-49.

Garavan H. Serial attention within working memory. (1998). *Memory & Cognition*. 26(2):263-76.

Gehring, W. J., & Knight, R. T. (2000). Prefrontal - cingulate interactions in action monitoring. *Nature Neuroscience, 3,* 516-520.

Jonides, J.,,Badre, D., Curtis, C., Thompson-Schill, S.L., and Smith, E.E. (in press). Mechanisms of conflict resolution in prefrontal cortex. In D.T. Stuss and R.T. Knight (Eds). *The Frontal Lobes*. Oxford: Oxford University Press.

Jonides, J., Reuter-Lorenz, P., Smith, E.E., Awh, E., Barnes, L., Drain, M., Glass, J., Lauber, E., Patalano, A., Schumacher, E.H. (1996). Verbal and spatial working memory, In D. Medin (Ed.), *The Psychology of Learning and Motivation*, 43-88.

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

Longoni, A. M., Richardson, J. T., & Aiello, A. (1993). Articulatory rehearsal and phonological storage in working memory. *Memory and Cognition*, *21*, 11-22.

MacDonald, A. W., Cohen, J. D., Stenger, V. A., and Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science*, 288, 1835-1838.

Meyer, D.E., and Kieras, D.E. (1997). A computational theory of executive processes and multiple-task performance: I. Basic mechanisms. *Psychological Review*, 104, 3-65.

Miller, E.K., Erickson, C.A. & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience 16*(16), 5154-5167.

Norman, D. A. & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior.(In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1-18). New York: Plenum.).

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

Salthouse, T. (1991). *Theoretical perspectives in cognitive aging*. Hillsdale, NJ.: Erlbaum

Schumacher, E.H., Lauber, E.J., Awh, E., Jonides, J., Smith, E.E., Koeppe, R.A. (1996). PET evidence for an amodal verbal working memory system. *Neuroimage*, *3*(2), 79-88.

Shallice, T. & Vallar, G. (1990). *Neuropsychological Impairments of Short Term Memory*. Cambridge: Cambridge University Press.

Smith, E.E., Jonides, J., & Koeppe, R.A. (1996). Dissociating Verbal and Spatial Working Memory Using PET. *Cerebral Cortex*, *6*(1), 11-20.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. In W. G. Koster (Ed.), *Attention and Performance II*, (pp. 276-31). Amsterdam: North Holland.

Sylvester, C-Y. C., Wager, T. D., Lacey, S.C., Jonides, J., Smith, E.E., and Nichols, T.E. (submitted). Comparing different executive mechanisms using neuroimaging: switching versus interference resolution. *Neuropsychologia.*

Figure Captions

Figure 1. A schematic of spatial and verbal working memory tasks and their respective control tasks. The figure show the events that occurred on typical trials. Note that the control tasks were designed to mimic as closely as possible the events in the memory tasks, with the storage requirement removed.

Figure 2. Brain activations for the verbal and spatial working memory tasks (shown in lighter greys for the higher levels of activation and in darker greys for the lower levels of activation) schematized in Figure 1. The top row of images shows brain activations of the memory-control subtraction superimposed on three views of a representative brain. The bottom row shows comparable activations for the spatial-control subtraction.

Figure 3. Brain activations superimposed on horizontal slices at the z-values indicated on the figure. Panel a shows the subtraction of the Search condition from the 2-back condition; panel b shows the subtraction of the Rehearsal condition from the 2-back condition. Activations are shown as lighter grays superimposed on darker greys that reveal the anatomy of a representative brain.

Figure 4. Panel a: A schematic of a task used to examine the hypothesis that spatial rehearsal recruits a process that influences the allocation of attention to visual objects. Panel b: Response times for probe items presented during the retention interval of the task shown in Panel a. The four bars represent whether the location matched or mismatched the location of the memorandum for the spatial and letter memory tasks respectively.

Figure 5. Panel a: A schematic of a task used to test whether allocation of attention to a visual location during a retention interval would affect memory for a spatial location that had to be stored. The retention interval was filled by either a large or small patch of color (represented by the striated region in the figure). Panel b: Accuracy of memory for the spatial location as a function of whether subjects had to perform a color discrimination during the retention interval (right pair of points) or merely viewed the colors but only had to perform the spatial memory task (left pair of points). Panel c: Accuracy on the color discrimination task when it was presented alone or with the spatial memory task.

Figure 6. Panel a: A schematic of a task used to study the brain basis of spatial working memory in posterior cortex. Subjects were sequentially presented with three characters to memorize, after which there ensued a 7 sec retention interval and a probe item. The memory task was either for the spatial location of the items, or, in another condition not shown, for the identities of letters presented at those locations. Panel b: Percent signal change in 6 posterior coronal slices for the spatial and verbal memory tasks.

Figure 7. Panel a. A task used to examine the effects of switching between representations of counters stored in working memory as well as inhibitory effects

induced by an incompatible response requirement compared to a compatible requirement. Panel b: A graph showing that in early stages of practice, there was a modest interaction in response times between the switching and compatibility variables, but with a small amount of practice, the interaction reduced in size sufficiently that there appears to be an additive relationship between switch-cost and the effect of compatibility.

Figure 8. Regions of common activation in the switching and compatibility tasks superimposed on a the canonical structural brain drawn from the Montreal Neurological Institute as included in the SPM program. Note that two views of the activations are shown to make it clearer where they are localized.

Figure 9. Panel a: Regions of greater activation in the switching than the compatibility task. Panel b: Regions of greater activation in the compatibility task than the switching task. For both panels, the activations have been superimposed on the canonical structural brain drawn from the Montreal Neurological Institute as included in the SPM program.