

tures damaged in amnesia (e.g., medial temporal lobe structures) play a role in forging these new bindings between preexisting elements in long-term memory.

The idea that amnesia is a deficit in long-term binding is appealing. Patients with amnesia exhibit severe deficits in episodic memory, or memory for experienced events. Fundamental to the definition of episodic memories is the idea that they are comprised of disparate temporal and contextual (external and internal) elements that uniquely determine each episode. If amnesic patients were unable to bind these elements in the long term, they would exhibit impaired retention of episodes once they are no longer held in working memory. However, amnesic patients, particularly those in whom medial temporal lobe damage extends outside of the hippocampus, are also impaired at learning new semantic information, such as facts and new vocabulary words (Squire & Zola 1998). They also exhibit impaired memory based on familiarity (Yonelinas et al. 1998). For these types of memory it is not clear how formation of semantic memories depends on long-term binding. Learning a new vocabulary word could be viewed as forming new connections between the phonological elements of the word and semantic elements. However, it is not obvious why amnesic patients would be impaired at this type of binding, but not the type of binding that supports non-declarative forms of learning that are intact in amnesia, such as priming for novel information, text-specific facilitation of reading speed with practice, and sequence-specific learning in the serial reaction time task (Squire et al. 1993). It may be that amnesic patients are specifically impaired in binding that leads to consciously accessible (or declarative) information. However, it seems somewhat more parsimonious to hypothesize that they are impaired at the formation of declarative memory representations, rather than in a particular set of binding processes.

In addition to studies of amnesia, equally important insights into memory function have been gained by research on how depth-of-processing manipulations facilitate later memory retrieval (Craik & Lockhart 1972). If one hypothesizes that the encoding of new memories results in the creation of new representations, one would suppose that deep, semantic encoding of information would create a richer, more interconnected, and more robust memory representation than when only surface features are processed. However, if one views the creation of new memories as the formation of new bindings between pre-existing elements, one would interpret deeper encoding as leading to more extensive binding. However, one might predict that if amnesic patients have deficient long-term binding processes, they would not benefit as much as neurologically intact individuals would from deeper processing at study. In fact, depth-of-processing effects are generally proportional for amnesic patients (Hamann & Squire 1996), suggesting that amnesia and depth of processing affect memory independently.

Another neuropsychological syndrome that is relevant to the study of memory representations is semantic dementia. Semantic dementia (SD) is a clinical term given to the temporal variant of frontotemporal dementia in which the affected patient shows progressive focal atrophy of the inferolateral aspect of the left and/or right temporal lobes with (in the early stages) relative sparing of the hippocampal complex (Graham & Hodges 1997). Although SD patients show a severe loss of semantic information (i.e., picture naming, verbal definitions of words, category fluency) and produce errors that reflect a loss of subordinate information while having superordinate information relatively preserved, they are able to recall recent life events with relative clarity and often perform relatively well on tests of nonverbal working memory (Hodges & Miller 2001). Patients with SD have also been found to show a temporal gradient in recognition of famous people that is the reverse of that shown by amnesic patients. Hodges and Graham (1998) reported that four patients with semantic dementia were better able to recognize currently famous individuals than previously famous individuals from a set of names of famous and nonfamous people. They were also better at identifying informa-

tion about currently famous people. These results provide compelling evidence that individuals with lateral temporal damage are able to store memories, initially, in a part of the brain (presumably the medial temporal lobes) that remains intact. Older knowledge, perhaps because it has been relocated to regions that are damaged by the disease process, does not fare as well. If long-term memories are simply bound representations of pre-existing knowledge, it is unclear why recent memories are spared if they depend on the same set of stored representations as more remote memories. It would seem that new memories and older memories should be equally affected by a loss of semantic knowledge. Rather, these data seem more consistent with the idea that recent memories depend on different neural representations than older memories. If memory representations depend on different neural substrates as they age, and if working memory is simply activation of long-term memory, this raises the question of how "activated" recent and remote memories may differ.

Ruchkin et al. have made a compelling case for their view based on electrophysiological and behavioral evidence from neurologically intact individuals. It is important to account for the data from neuropsychological patients, as well. It may be that their model can be made consistent with neuropsychological findings, if certain assumptions are made as to the processes impaired in these patients. Hopefully, these assumptions will be testable as a means to further assess their model.

ACKNOWLEDGMENTS

This work is supported by a National Science Foundation CAREER Award (IBN-9985417) (B.J.K.) and by a Julie-Payette NSERC Research Scholarship (I.V.V.).

Long-term memories, features, and novelty

James K. Kroger

Department of Psychology, New Mexico State University, Las Cruces, NM 88003. jimkk@umich.edu
<http://psych.nmsu.edu/~jkroger/index.html>

Abstract: Ruchkin et al. make a strong claim about the neural substrates of active information. Some qualifications on that conclusion are: (1) Long-term memories and neural substrates activated for perception of information are not the same thing; (2) humans are capable of retaining novel information in working memory, which is not long-term memory; (3) the content of working memory, a dynamically bound representation, is a quantity above and beyond the long-term memories activated, or the activity in perceptual substrates.

Ruchkin et al. present compelling evidence that information in working memory, rather than existing in a special purpose buffer distinct from the neural substrates specialized for perceiving that kind of information, is a state of activation in those same substrates under the control of frontal cortex. As the authors note, this is a more parsimonious scheme than duplicate representation architectures for the perception and storage of the myriad kinds of information we deal with. The view that attention activates representations, even in low-level visual areas, has also been demonstrated for nonverbal information by Kastner et al. (1999) and others, and the control of posterior representations by frontal cortex was embodied in our computational model of working memory (Kroger et al. 1998).

It might be possible to make a more explicit distinction between "long-term memory" in the sense of semantic information that is activated (e.g., in inferior temporal cortex) during retention of verbal information in working memory, and the neural substrates that are engaged as information is perceived. In terms of verbal information, there may be a great deal of overlap, as perceived words typically activate areas specific to semantic information storage. However, as a general principle of information storage, some caution is called for. It is well established that cells in primary visual

cortex participate in perception of light and dark, orientation, and so on. Although it seems straightforward that retaining information in working memory about a fixation dot, or the location of a saccade target, may involve activation of regions that participated in its perception, it is less clear that this is an activated long-term memory. Suppose subjects are presented with a novel shape to remember. It could be argued that elements of the shape (corners, curves, etc.) are retained by invoking neural representations of these features learned over time, thus perhaps constituting long-term memories bound in the current episode into the novel shape. In the case of five novel shapes, of different colors, the argument could be similarly made that the content of working memory is a binding of long-term memories. It seems, however, that there is a meaningful sense in which the “content” of working memory is more a binding of features to constitute a new representation than it is activated long-term memories. In the same vein, the proposition “John loves Mary” is distinct from the proposition “Mary loves John,” although the same long-term memory elements may be activated by both. What if one has never met John, does not know which John is being referred to, or has never encountered the name “John” before? The degree to which the content of working memory includes activated long-term memory varies, and semantic and episodic associations activated along with the proposition may vary, but the simple proposition “John loves Mary” can exist in working memory apart from this extra information. As with novel shapes, it may be argued that novel information is retained that does not depend on long-term memories.

The relocation of working memory content from separate buffers, as proposed by Baddeley (1986), to the neural substrates specialized for perception of information is an important and necessary step that will help enable the study of how frontal cortex and attention accomplish that which Ruchkin and his coauthors term “episodic” bindings. Now that neuroimaging methods permit observation of the actual substrates of memory retention, the notion that separate regions exist for buffering and for perception of information could interfere with proper interpretations of neuroimaging results, and is no longer tenable. This is especially true because these buffers were often associated with frontal activity observed during retention, which more likely is involved in attentional control of working memory, as the authors suggest. Long-term memory may become activated as meaningful stimuli, such as words, are retained in working memory, but it is important to remember that humans are facile at retaining novel information, and novel, complex bindings of information. In these instances, the nature of “what is in working memory” transcends the collection of activated perceptual substrates or long-term memories that may be active, especially in the context of research on higher cognition.

Some neuroimaging results suggest that, as the representations bound together in working memory increase in complexity, but not as memory load increases, more anterior regions of prefrontal cortex are recruited (Kroger et al. 2002).

Working memory as a mental workspace: Why activated long-term memory is not enough

Robert H. Logie and Sergio Della Sala

Department of Psychology, University of Edinburgh, Edinburgh EH8 9JZ
Scotland, United Kingdom. r.logie@staffmail.ed.ac.uk
sergio@abdn.ac.uk

Abstract: Working-memory retention as activated long-term memory fails to capture orchestrated processing and storage, the hallmark of the concept of working memory. The event-related potential (ERP) data are compatible with working memory as a mental workspace that holds and manipulates information on line, which is distinct from long-term memory, and deals with the products of activated traces from stored knowledge.

Ruchkin et al. equate working-memory retention systems with short-term memory. One advantage of the multiple-component working-memory concept is that it incorporates both on-line processing and temporary memory (e.g., Baddeley & Logie 1999), allowing the concept of working memory as a mental workspace (Logie 1995), rather than as a simple temporary storage device. The notion of working memory retention systems as comprising a state of activated long-term memory fails to capture, or to account for, this concept of orchestrated processing plus storage. A multiple-component working memory, as a mental workspace that is separate from, but holds and manipulates the products of activated traces in long-term memory (Logie 1995; 2003), retains the advantages of offering a testable theory, while accounting for a wide range of behavioural data, both from experimental manipulations and from neuropsychological dissociations (Della Sala & Logie 2002).

Ruchkin et al. argue that it is more parsimonious to assume that short-term memory reflects simply the activation of long-term memory traces, than to assume a separate, multiple-component working-memory system. However, the former theory has to make a wide range of assumptions regarding: the operation of the activation process, how thresholds are set or adjusted, how non-relevant but activated traces are inhibited, how the temporary binding process occurs and is maintained with input from different modalities, how the activated information is manipulated (anything from backward digit recall to generating novel mental images), and how the novel results of those manipulations are held on a temporary basis (for detailed discussion, see Logie 2003). Moreover, a model that explains dual-task interference in terms of similarity of the codes used for each task sounds dangerously circular. Suggesting that two tasks interfere because they use similar codes has some difficulty in making predictions independently of the experimental outcome (Cocchini et al. 2002). In other words, assuming that working-memory retention systems and long-term memory arise from the same conceptual cognitive systems may well be theoretically sterile.

One feature of experimental research into human cognition that is all too rarely recognised is the use of a range of cognitive strategies by participants. Logie et al. (1996) demonstrated that even very simple immediate serial-ordered recall tasks are prone to the use of a range of cognitive strategies, both across individuals, and within the same individual from one occasion to another, even if the aggregate data for the sample of participants generate reliable behavioural phenomena. Indeed, some fMRI data of our own (Logie et al., in press) have shown that specifically instructing participants to use subvocal rehearsal results in a relatively narrow range of areas of activation, focused on Broca's area and the supramarginal gyrus, compared with the much broader network of activation patterns, including those same areas, reported by Paulsen et al. (1993). Interpreting precisely what kind of cognitive function might be employed for any given task requires a very careful cognitive-task analysis, with independent behavioural evidence to indicate precisely how participants are performing the tasks. Without this, at best, the activation patterns can be seen as correlates of how, on average, the participant group performed the task they were set. This is not necessarily informative about the cognitive systems that participants may select to meet the requirements of the task in hand. This kind of detailed cognitive-task analysis is not common in brain imaging studies, and this makes it very difficult to suggest that a given aggregate pattern of activation is reflecting any particular cognitive function. The result may then be a mapping of tasks onto brain structures rather than a mapping of cognitive functions onto brain organisation. In this sense, brain activation patterns might be used to confirm a cognitive theory, but they need not constrain that theory.

The bulk of the evidence described by Ruchkin et al. focuses on the argument that activation of the same brain areas indicates that the same cognitive function is involved. Specifically, because the same brain areas are active for tasks that are assumed to require temporary retention as are involved in activation of stored knowl-