ideas from connected prose. The problem for these patients appears to be in their ability to consolidate new information into LTM. Thus, they form temporary representations that use LTM, but these representations cannot be stored in LTM. Considerations such as these suggest that the unitary-system view does not necessarily provide a more plausible account of available literature than does the multiple-component view.

A truly plausible account of the available literature must be able to explain fundamental and well-established laboratory phenomena. For example, the word-length effect, wherein fewer long words can be retained over a delay interval than short words, suggests that these STM representations are minimally distinguished on the basis of their meaning and are principally phonologically-based (Baddelely et al. 1975). If, as the unitary account claims, information in the focus of attention is activated from long-term memory, one would expect the activated representations to be maximally distinguishable based on the rich and semantically elaborate nature of long-term representations. Similarly, the semantic nature of LTM representations would seem to inoculate them to the confusability created by phonologically similar items in to-be-retained lists. Yet, such confusability is reliably observed in the laboratory (Baddeley 1966a; 1966b; Baddeley & Dale 1966). A possible clue to explain the emergence of these effects is the observation of phonological similarity and word-length effects in sign language (Ronnberg et al. 2000; Wilson 2001). This differs from the typical description of the sketchpad in nonsigners, which indicates that sequential information is not well-maintained by the visuospatial sketchpad (Baddeley 2001a).

An explanation of capacity limits has been considered central to completely understanding the architecture of the human mind. The unitary system provides one plausible account, based on limitations of the focus of attention. However, the unitary-system account cannot easily explain observed capacity-limit differences that depend on the composition of to-be-remembered lists. That is, immediate memory for a list of randomly ordered words averages between 5–7 words, but memory for words in sentences averages between 13–22 words (Craik & Masani 1969). Additionally, it has been demonstrated that capacity limits can be overcome by training individuals to chunk information (Ericsson et al. 1980). The observation of multiple capacity limits implies the operation of multiple memory systems. Chunking in retrieval structures apparently can be used to overcome capacity limits, but their role in a unitary memory system is not clear (Ericsson & Kintsch 1995; Gobet et al. 2001).

Ruchkin et al. argue that the superior temporal resolution offered by ERP can yield evidence to distinguish between unitary- and multiple-systems views. Although it is indeed true that ERP does offer superior temporal resolution, many of the changes in the data they show occur on the order of four seconds or more, a time scale certainly resolvable by fMRI (Zarahn et al. 1997). Moreover, more precise evidence for the claims made by Ruchkin et al. may come from techniques with superior spatial resolution, such as fMRI. Under these circumstances then, ERP and fMRI should converge, and, in fact, they do. In a number of studies, for example, PFC activation similar to that shown by Ruchkin et al. has been observed (Rypma & D'Esposito 1999; 2000; 2001; Rypma et al. 1999). These studies complement the results shown by Ruchkin et al. to the extent that they show activation increases in PFC following onset of to-be-remembered information. Further, they show distinct activation topographies based on whether to-be-remembered lists were short (i.e., 2–3 letters) or long (i.e., 6 letters). That is, ventral regions of PFC showed activation during retention of both short and long lists. Dorsal PFC, however, showed a more load-sensitive activation pattern. That is, there was minimal activation during retention of the short list, but substantially increased activation during retention of the long list.

In summary, although these results do not provide a critical test between the multiple and unitary working-memory theories, they do support the notion that multiple cortical regions are involved in STM maintenance, depending on task demands. We interpret these results to indicate the existence of separate STM systems used to support information retention under high-memory-demand conditions.

The short-term dynamics within a network of connections is creative

William A. Phillips

Department of Psychology, Centre for Cognitive and Computational Neuroscience, University of Stirling, Stirling FK9 4LA, United Kingdom.
wap1@stir.ac.uk

Abstract: Although visual long-term memory (VLTM) and visual short-term memory (VSTM) can be distinguished from each other (and from visual sensory storage [SS]), they are embodied within the same modality-specific brain regions, but in very different ways: VLTM as patterns of connectivity and VSTM as patterns of activity. Perception and VSTM do not "activate" VLTM. They use VLTM to create novel patterns of activity relevant to novel circumstances.

There is strong evidence for several components of Ruchkin et al.'s thesis, including much that is stronger than that which they cite. I was surprised to see a review of the role of prefrontal cortex (PFC) in working memory (WM) that neglects the work of Pat Goldman-Rakic. Furthermore, although Ruchkin et al. distinguish between different components of visual memory, they make no use of the first and strongest evidence distinguishing SS, VSTM, and VLTM (Phillips 1974; 1983b; Phillips & Christie 1977a), nor do they relate them to general purpose processing capabilities, such as those thought to involve PFC (Phillips & Christie 1977b). The possibility that temporary dynamic storage involves activity within a network in which LTM is embodied in the connection strengths is an idea of great antiquity, for which there is ample evidence. Phillips (1983b, Fig. 1) presents just one of many versions of this idea as applied to visual memory. This does not imply that VSTM occurs within *the same* systems as those that are initially processing the information, however. The regions within which SS occurs are included in the latter, but not the former (i.e., that in which VSTM occurs). If activity in striate and peristriate areas could be maintained, then SS (photographic or iconic memory) would be a voluntary option. It is not (Phillips 1974; 1983b; Phillips & Singer 1974; Simons & Levin 1997). Thus, WM is possible within some cortical regions, but not others. The evidence reviewed by Ruchkin et al. suggests that voluntary maintenance may not be possible in regions of the visual stream prior to its division into dorsal and ventral pathways.

The dependence of short-term dynamics on long-term changes in connectivity are so important for our understanding of cognition that it is necessary to ask whether the notion of "activation" clarifies this relationship. I don't think that it does. First, activity and connectivity are very different things. Neural network studies clearly show that the short-term dynamics of a network with recurrent connections can be very complex, and evolve on a short-term time scale into many different patterns of activity, without any changes in the connection strengths. Second, Ruchkin et al. use the word "activation" in several different ways, and their equivalence is far from obvious. Sometimes they use it in a psychological sense, sometimes to refer to EEG measures, and sometimes to refer to underlying neuronal activity. The relation between the latter two is distant, and that between both of these and the former is even more distant. For example, the ordinate in their Figure 3 refers to "activation," but it doesn't seem to really mean measured EEG activity, as in some of their other figures. It is even harder to relate Figure 3 to neuronal activity, as the activity of neurons in the visual cortex has a far more complex time-course in response to stimulation than that given in Figure 3 as a representa-

https://doi.org/10.1017/S0140525X03480163 Published online by Cambridge University Press

tion of iconic activity. For example, contrast Ruchkin et al.'s Figure 3 with Figure 1 in Phillips (1983a), which is based upon the single-unit neuronal activity reported in Singer and Phillips (1974). Third, "activation" connotes the selection of old things, not the creation of new things. Even pre-attentive perceptual processes must involve the latter (Watt & Phillips 2000).

There is space here to outline only two paradigms in support of my view that the creation of new descriptions cannot usefully be viewed as simply the activation of old items. The first involves the use of matrix block patterns to compare VSTM and VLTM (Phillips 1983b; Phillips & Christie 1977a). To focus on visual memory, these patterns were designed to be difficult to verbalize adequately. To focus on the use of novel descriptions, rather than on the use of old items, novel patterns were used on every trial. The results clearly showed that accurate descriptions of novel patterns could be voluntarily maintained for as long as the subject could keep attending fully to them, providing the matrix patterns were below measured levels of complexity. Some VLTM for those patterns was produced by a single presentation, but with much less accuracy than for VSTM. I cannot see how it is useful to think of these VSTM descriptions as activation of long-term memory items, unless by "activation," Ruchkin et al. mean the use of descriptive capabilities stored in VLTM to create novel descriptions.

The second paradigm involves mental rotation. To see whether rotated images formed within VSTM operate upon the long-term representations of the items rotated, we modified a paradigm that studies discrimination between normal and mirrored versions of familiar alphanumeric characters displayed in various orientations (Shepard 1978). Our results show that when subjects image a familiar alphanumeric character upside down they still have the upright character in VLTM available for use (Roldan & Phillips 1980). This shows that mental rotation uses knowledge of the form of the upright character and transformational rules to create a novel description within VSTM. They do not transform the VLTM of the familiar form.

I assume that there will be general agreement with Ruchkin et al.'s view that many modality-specific brain regions embody both VSTM and VLTM. I do not see how the concept of "activation" advances our understanding of these issues, however. That said, my reservations concerning the conceptual framework within which Ruchkin et al. interpret their data do not imply that measures of the kind they emphasize are without value. Many, such as those of EEG coherence, may well be of value, if validly related to the underlying neuronal activity.

Models versus descriptions: Real differences and language differences

Jeroen G. W. Raaijmakers[a] and Richard M. Shiffrin[b]

[a]*Department of Psychology, Universiteit van Amsterdam, Amsterdam, 1018 WB The Netherlands;* [b]*Cognitive Science Program, Indiana University, Bloomington, IN 47405.* **J.G.W.Raaijmakers@uva.nl shiffrin@indiana.edu**

Abstract: We argue that an approach that treats short-term memory as activated long-term memory is not inherently in conflict with information recycling in a limited-capacity or working-memory store, or with long-term storage based on the processing in such a store. Language differences aside, real model differences can only be assessed when the contrasting models are formulated precisely.

The authors of the present target article introduce their proposal in an attempt to draw a distinction between the idea of short-term or working memory as a separate store and the idea of short-term or working memory as activated representations in long-term memory. We note that these views are not inherently opposed to one another. In particular, it makes a great deal of sense to assume that short-term memory corresponds to activated representations in long-term memory, but that does not mean that a model for short-term memory that is based on the notion of stores or buffers is useless or even incorrect.

In this respect, it is of some interest to note that the Short-Term Store (STS) system in the modal model of memory proposed by Atkinson and Shiffrin (1968) has long been formulated as the temporarily activated portion of Long-Term Store (LTS) (e.g., Atkinson & Shiffrin 1971). This reformulation was based on theoretical grounds: in particular, the idea that it made more sense to assume that perceptual stimuli contact information in long-term memory, rather than to assume a sequence from sensory registers to short-term store to long-term store. This idea was further elaborated in Shiffrin (1975; 1976). Of course, the idea of short-term memory as activated representations in long-term memory considerably predates Atkinson and Shiffrin, going back, at least, to James (e.g., 1890).

More importantly, this idea is easily reconciled with a model that assumes that STS or working memory may be viewed as a store that temporarily holds a small amount of information for further (more elaborate) processing. As argued by Shiffrin (1975; 1976), perceptual information activates a large amount of long-term memory information. However, the information is rapidly lost from STS (i.e., becomes inactive) unless it is maintained in STS through rehearsal and other coding processes. As a result, only a few items may be maintained simultaneously in a highly active state in STS. A STS buffer, such as that proposed by Atkinson and Shiffrin (1968), is a simple model used to describe this process of maintenance of information in STS.

The history of the psychology of memory has shown a number of examples where ideas that are not necessarily mutually exclusive lead to unfruitful debates. Perhaps the clearest example is the way in which the Atkinson–Shiffrin modal model of memory is usually discussed in textbooks and put into opposition with the levels-of-processing framework (Craik & Lockhart 1972) or the working-memory model of Baddeley and Hitch (1974; 1977). The target article echoes these textbook accounts when it mentions that "the modal model . . . does not provide an accurate account of how short-term and long-term memories interact, nor does it correctly predict performance for certain dual-task experiments" (sect. 1.1). However, Raaijmakers (1993) and others (Bjork 1975; Glanzer 1977; Shiffrin 1977) have argued that the conflict between the levels-of-processing approach and the Atkinson–Shiffrin model is artificial and not based on a detailed analysis of the Atkinson–Shiffrin model. In particular, it does not take into account the role assigned to the control processes of rehearsal and coding (or maintenance and elaborative rehearsal). In a similar vein, it has been argued that the evidence that was put forward by Baddeley and Hitch (1974) does not really contradict the modal model (see Raaijmakers 1993).

We are afraid that the present target article might similarly promulgate a false dichotomy and help initiate a flood of papers showing either the fruitfulness of the "store"-approach or arguing for the temporary-activation approach. Although some might see such a state of affairs as a sign of healthy progress, active debate is not always a good thing when the debaters are talking "past one another." We believe that such theoretical controversies are best resolved by careful and precise formulation of the different approaches, so that the fundamental and underlying similarities and differences can be assessed. Our personal approach has been to produce such specification by formulating mathematical and computer-simulation models. We predict that such formulations would show that the two model types are not in conflict, but rather that each type has many differing variants that would be amenable to experimental testing. In addition, the choice of model representation may be more a matter of style than substance. The preferred choice of "stores" or "activated subset" will probably depend most on the nature of the data that one tries to accommodate, and an assessment of which approach proves more fruitful, parsimonious, or productive. As such, the situation is reminiscent of the wave versus particle viewpoints in contemporary physics.