Design for a Working Memory

Oberauer, Klaus

Abstract: Starting from the premise that working memory is a system for providing access to representations for complex cognition, six requirements for a working memory system are delineated: (1) maintaining structural representations by dynamic bindings, (2) manipulating structural representations, (3) flexible reconfiguration, (4) partial decoupling from long-term memory, (5) controlled retrieval from long-term memory, and (6) encoding of new structures into longterm memory. The chapter proposes an architecture for a system that meets these requirements. The working memory system consists of a declarative and a procedural part, each of which has three embedded components: the activated part of long-term memory, a component for creating new structural representations by dynamic bindings (the ''region of direct access'' for declarative working memory, and the ''bridge'' for procedural working memory), and a mechanism for selecting a single element (''focus of attention'' for declarative working memory, and ''response focus'' for procedural working memory). The architecture affords two modes of information processing, an analytical and an associative mode. This distinction provides a theoretically founded formulation of a dual-process theory of reasoning.

CHAPTER TWO

DESIGN FOR A WORKING MEMORY

Klaus Oberauer

Contents

Abstract

Starting from the premise that working memory is a system for providing access to representations for complex cognition, six requirements for a working memory system are delineated: (1) maintaining structural representations by dynamic bindings, (2) manipulating structural representations, (3) flexible reconfiguration, (4) partial decoupling from long-term memory, (5) controlled retrieval from long-term memory, and (6) encoding of new structures into long-term memory. The chapter proposes an architecture for a system that meets these requirements. The working memory system consists of a declarative and a procedural part, each of which has three embedded components: the activated part of long-term memory, a component for creating new structural representations by dynamic bindings (the "region of direct access" for declarative working memory, and the "bridge" for procedural working memory), and a

mechanism for selecting a single element ("focus of attention" for declarative working memory, and "response focus" for procedural working memory). The architecture affords two modes of information processing, an analytical and an associative mode. This distinction provides a theoretically founded formulation of a dual-process theory of reasoning.

1. SIX REQUIREMENTS FOR A WORKING MEMORY SYSTEM

The study of working memory (WM) has its roots in the investigation of immediate recall of short lists of items such as a telephone number or a list of unconnected words (Blankenship, 1938; Nipher, 1878). Until today most research on WM is devoted to what has become known as *memory span* tasks. The seventies of the last century have witnessed an extension of the scope of the concept of WM, now defined as a device for simultaneous storage and processing (Baddeley & Hitch, 1974), and with it came the introduction of so-called *complex span* tasks that combine the immediate recall of lists with a concurrent processing task such as reading (Daneman & Carpenter, 1980). The construct such developed has become a big success not least because complex span tasks have turned out to be strong predictors of performance in complex cognitive activities such as text comprehension and reasoning (Daneman & Merikle, 1996; Kyllonen & Christal, 1990; Süß, Oberauer, Wittmann, Wilhelm, & Schulze, 2002), and are strongly related even to general intelligence (Conway, Kane, & Engle, 2003).

As theorists have recognized for some time now (e.g., Baddeley, 1986; Just & Carpenter, 1992), WM has certainly not evolved to help us remember telephone numbers (whether on their own or while reading unrelated sentences). This insight nourishes the suspicion that the tasks commonly used to operationalize WM are more a reflection of the historical development of the field than of our theoretical understanding of its object. I think this is unfortunate because much research is devoted to understanding the particular tasks that happen to have emerged as established operationalizations of WM, asking questions such as "What happens in the reading span task?" instead of, "How does working memory work?"

In this chapter, I want to pursue a functional approach to WM. I will start from an assumption about what WM is good for and attempt to delineate some minimal requirements of a system that serves this function. This approach is akin to Marr's (1982) strategy to understand vision through a computational analysis of what the visual system must accomplish, although my analysis will not yet reach the degree of formal precision that he achieved. The leading idea is that form follows function, so that an analysis of the function a system is to accomplish provides useful constraints for its architecture and its mechanisms.

I assume that WM is a system devoted to providing access to representations for goal-directed processing. At least in humans this involves processes such as language comprehension, reasoning, planning, hypothetical thinking, and creative problem solving. Success in all these activities, which I will refer to as complex cognition, is highly correlated with measures of WM capacity (for a review see Oberauer, Süß, Wilhelm, & Sander, 2007). One common characteristic of all varieties of complex cognition is that they require the flexible combination of representations into new structures (Halford, Wilson, & Phillips, 1998; Oberauer et al., 2007), and the goal-directed manipulation of these structures.

I propose that a system serving complex cognition must meet at least the following six demands: First, it must be able to build and maintain new structural representations. For instance, the system must represent new propositions, new sequences of actions in a plan, or new constellations of pieces on a chessboard. Building new structural representations requires a mechanism for *dynamic binding*, that is, for binding content elements such as words, objects, or events, to places in a cognitive coordinate system or to variables in a schema. For instance, objects must be bound to locations in physical space, events to locations on the temporal dimension, and words to variables in syntactical and propositional schemata. These bindings must be dynamic, which means that they can be set up quickly and dissolved quickly when the structure is updated or discarded.

Second, complex cognition involves manipulating structural representations. This means that there must be a mechanism for selectively accessing one or a few elements within a structure and submitting them to a cognitive operation. For instance, mentally simulating a chess move involves picking out one figure and moving it across the board, while all other figures stay in their positions. This requirement calls for an *attentional selection* mechanism for contents of WM. In addition, the system must have a mechanism for deciding what to do with the selected element. This requirement points to the need for a *procedural* system that represents condition–action rules (e.g., procedures for deciding which moves to consider in a mental simulation, given a constellation on the chessboard, and procedures for moving the figures).

Third, WM is a general-purpose mechanism, not a module designed for solving a particular problem. Therefore, it must be possible to flexibly reconfigure it. This requirement points to the need for *executive processes*, that is, processes that control its operations and configures its parameters according to representations of goals. For instance, WM would operate differently when the goal is to remember a constellation of pieces on a chessboard than when the goal is to play a game, and it would operate differently in speed chess than in ordinary chess.

The fourth requirement follows from the need for rapid updating of structural representations and of the procedures acting on them:

Representations in WM must be partially decoupled from knowledge and episodic memories in *long-term memory* (LTM). A chess player mentally simulating two or three moves must be able to distinguish the representation of a hypothetical board constellation from the many representations of similar constellations that he remembers from previous games. He must also be able to distinguish the constellation that follows after an imagined move from the constellation before that move. In general, manipulation of structural representations means that the structure is frequently updated, and therefore, WM must be able to avoid proactive interference from the previously held representations on the current one. Likewise, the flexible implementation of different goals requires frequent updating of procedural representations, and therefore, WM must also be able to avoid proactive interference from habits and routines.

Fifth, WM needs to be able to draw on relevant contents of LTM. This implies that LTM representations that could be helpful for the problem at hand can be retrieved efficiently. This need creates an obvious tension with the requirement for decoupling WM representations from LTM. Thus, the system must decide on when to allow LTM to influence the contents of WM and when to block them off.

Finally, working out the solution to a problem would lose much of its worth if the solution would be quickly forgotten once the system is engaged with something else. Therefore, new structural representations built in WM must be transferred into LTM. This means that there must be a mechanism for transforming temporary, dynamic bindings into more permanent structural representations.

In this chapter, I will propose a blueprint for a system that meets these requirements. Section 2 will sketch the architecture of WM, delineating mechanisms for meeting the first three requirements. Section 3 will discuss the interaction between WM and LTM, offering some speculations on how the system meets the remaining three requirements. The structure of the chapter does not follow strictly the list of the six requirements; therefore, Table 1 makes explicit which component of the WM system address which requirement and provides pointers to the sections explaining these components.

2. THE ARCHITECTURE OF WORKING MEMORY

I will propose a sketch of the architecture of WM that is suited to fulfill the functions introduced above. The core of the framework has first been developed to accommodate empirical evidence (Cowan, 1988; Oberauer, 2002); here I will flesh it out by giving it a functional interpretation, and by adding further assumptions and components, hoping that the elaborated framework will prove useful as a blueprint for a system that serves complex cognition.

Table 1 The Six Requirements for a Working Memory System, the Components Assumed to Be Responsible for Them, and Where They Are Addressed in the Text.

Requirement	Component	Sections
1. Structural representations	Region of direct access	2.1.2
	Bridge	2.2
2. Manipulation	Focus of attention	2.1.3
	Procedural WM	2.2
3. Flexible reconfiguration	Executive processes	2.2.2
4. Partial decoupling from LTM	Distinction between activated LTM and central components	2.1
	Control of retrieval from LTM	3.3
5. Retrieval from LTM	Retrieval and unpacking of chunks	3.2.2
6. Encoding of structural information into LTM	Chunking of structural information	3.2.1

A first and fundamental design decision is to distinguish between declarative and procedural WM. The declarative part is responsible for making representations available for processing, whereas the procedural part is responsible for doing the processing. In a sense, the declarative system is the memory part, and the procedural system is the working part of working memory. The distinction between declarative and procedural memory lies at the heart of production system architectures such as ACT (Anderson, 1983), SOAR (Newell, 1990), and ACT-R (Anderson & Lebiere, 1998), but its implications for WM have not been spelled out (for one previous reference to "procedural working memory" see Monsell, 2003). In much of the WM literature its place has been taken by the distinction between the *central executive* and the *slave systems* (Baddeley, 1986). In my view this is an unfortunate conceptualization because the procedural part of the system is no more central than the declarative part. Others have treated the terms *working memory* and *executive system* as largely equivalent, or the former as a subset of the latter. This is even more unfortunate because it blurs the important distinction between the declarative and the procedural aspects. Moreover, not every process executed can be regarded as an executive process without draining the concept *executive* of all useful meaning. Here I will make a distinction between the primary processes carrying out the manipulations required for a task, such as mentally moving a pawn on a chessboard, or deciding whether to press the left or the right button in a speeded choice task, and executive processes that supervise and control the primary processes, such as the decision to switch to another task, to update the current WM contents by new information, or to sacrifice accuracy for speed in a choice task.

2.1. Declarative Working Memory

The declarative part of WM, illustrated in Figure 1, consists of three components that constitute three embedded sets of representations, the activated part of LTM, the region of direct access (DA region), and the focus of attention. Functionally, the three components can be seen as three levels of selection of representations, with each component narrowing down the selected set more than the preceding one. In addition, the three components have qualitatively distinct functions for the construction and manipulation of representations, as I will elaborate below.

2.1.1. Activated Long-Term Memory

I conceptualize LTM as an associative network of representations that activate each other automatically along their associations. The advantage of such a system is that it is content addressable: Given any piece of information, related information can quickly be activated, such that a
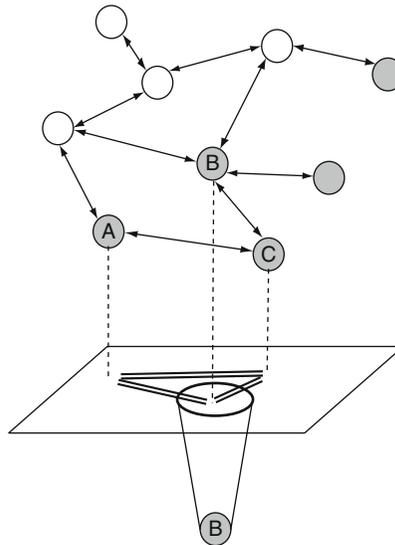
Figure 1 *Architecture of declarative working memory.* Small circles represent elements of declarative representations in long-term memory. Shaded units are activated above baseline. A subset of three elements (labeled A, B, and C) are bound to positions in a mental space, depicted here as the rectangular frame. This subset of elements constitutes the content of the region of direct access. Thin continuous arrows are associations in LTM; dotted lines are temporary bindings; double lines are relations between elements in the region of direct access that emerge from their relative positions in the mental space. One of the positions in mental space is selected by the focus of attention (large thick–lined oval); the element bound to that position (B) is thereby selected for processing.

whole set of tightly interconnected representations can be recreated from any subset as cue. This idea is incorporated in virtually all current theories and models of LTM (e.g., Anderson & Lebiere, 1998; Gillund & Shiffrin, 1984; McClelland, McNaughton, & O'Reilly, 1995).

An assumption that is not shared by all models is that, in LTM, there is no strict separation between declarative and procedural knowledge. Rather, declarative representations are associated with procedures, such that the activation of a declarative representation such as an object, an event, a situation, or a fact automatically activates an associated cognitive operation (e.g., drawing an inference based on an activated stereotype), evaluation (e.g., assessing an event as positive), or physical action (e.g., performing a left-directed action in response to a stimulus on the left side). These associations enable information processing that bypasses the more central components of WM. I will return to the power and the limitations of associative information processing in Section 4.

Representations in LTM are activated by perceptual input or through spread of activation from other, associated representations, including representations of goals. The currently activated subset of representations forms the activated part of LTM. The degree of activation of a representation is an implicit code of its expected relevance for the current situation and the current goal (Anderson & Lebiere, 1998), and its function is to make more relevant representations more available for processing.

The increased availability of activated representations is reflected in three consequences of activation. First, perceptual stimuli corresponding to already activated representations in LTM are processed more efficiently and can thereby be identified (i.e., categorized) faster and with higher accuracy. Thus, activation contributes to priming. This gain in perceptual fluency generates a signal of familiarity of the stimulus that can be used, among other things, to inform recognition decisions (Whittlesea, Jacoby, & Girard, 1990).[1]

Second, a representation is more likely to be retrieved, and its retrieval takes less time, the higher its activation (for a formal development of these assumptions see Anderson & Lebiere, 1998). Retrieval means to project a representation into the more central component, the region of direct access, where it can be manipulated and combined with other representations. The third consequence is that activated declarative representations, through their associations to procedural representations, generate tendencies for or against associated inferences and decisions, a point to which I will return Section 4.

---

[1] It should be added that both repetition priming and familiarity cannot be attributed only to temporary activation of representations, because both can last for times much longer than activation can plausibly be assumed to continue—in one instance, up to 17 years (Mitchell, 2006). I assume that activation in LTM drives familiarity only in short-term recognition paradigms in which the recognition decision follows initial presentation within a few seconds.

2.1.2. The Direct-Access Region

The main function of the DA region is to render a small number of distinct elements immediately accessible, and to integrate them into a structure. There are two ways of building new structures, both of which involve binding of content representations to context representations. One is to draw on existing structure templates or schemata, such as the schema for a proposition, which has an action at its core and relates it to arguments such as the role of an agent (''who does it?''), an object (''to who or what is it done?''), and other, optional roles such as instrument (''with what was it done?'') or time (''when did it happen?''). The template arguments are variables to which elements from large classes of concepts can be bound, enabling us to build a virtually infinite variety of structures. For instance, we can make sense of linguistic input such as: ''The parrot beats the sheep with a cucumber'' by binding the concept of *beating* to the action variable, the concept of *parrot* to the agent variable, and so on.

Another way to build structures is by binding each element to a position in a common cognitive coordinate system. A cognitive coordinate system is a mental space that can be used to literally represent physical space, or to metaphorically represent other continuous dimensions such as time or some quantitative variable (e.g., size, brightness, pitch, intelligence, etc.). For example, I could represent my knowledge that the population of China is larger than that of India by placing a token or symbol for China on top of one for India in my mental space. I could then proceed to add the knowledge that the population of India is larger than that of Russia by placing a representation of Russia below that of India. By being placed in a common coordinate system every element is related to every other, and this can bring out new relations that were not explicitly represented before, for instance, the relation between the population sizes of China and Russia. This relation is constrained by the spatial nature of the coordinate system that, in this case, enforces that Russia is placed below China, so that I can infer that the population of Russia must be smaller than that of China even if I never thought about this relation before. This is, in a nutshell, the mechanism of deductive reasoning as described by the theory of mental models (Goodwin & Johnson–Laird, 2005; Johnson–Laird & Byrne, 1991).

Binding elements into a multidimensional cognitive coordinate system is also indispensable for inductive reasoning. The core operation of inductive reasoning is to compare two entities to identify similarities and differences between them. To go beyond global judgments of similarity, the reasoner must approach the task in an analytic way, that is, distinguish relevant features of the two entities and make separate comparisons with regard to each feature. For instance, a person can judge that an orange and a tennis ball are similar with regard to size and color, but differ with regard to edibility. These relations can be represented simultaneously by placing the two

objects into a three-dimensional coordinate system in which the three dimensions are assigned to the three feature dimensions selected for comparison. Simultaneous representation of similarities and differences on several feature dimensions is a key in many inductive-reasoning tasks that feature in intelligence tests, such as the Raven matrices (Klauer, 1996).

The examples discussed above illustrate that the DA region recruits a spatial medium of representation as a projection screen for relations on nonspatial dimensions, by assigning to its dimensions, for instance, physical space, time, and any feature dimension on which entities could be compared. The prevalence of spatial metaphors for nonspatial relations has been noted in studies reaching from semantics (Lakoff & Johnson, 1980) to logical reasoning (DeSoto, London, & Handel, 1965; Knauff, Mulack, Kassubek, Salih, & Greenlee, 2002) to social psychology (Bar-Anan, Liberman, Trope, & Algom, 2007). I believe this prevalence reflects the fact that space is used as a generic medium for building structural representations in the region of direct access.

To summarize, the region of direct access is a mechanism for establishing and holding temporary bindings between contents (e.g., objects, events, words) and contexts (i.e., argument variables in structure templates, or positions in a generic cognitive coordinate system). By supporting arbitrary bindings between virtually any content with any context, this system enables the compositionality of thought that many theorists regard as a hallmark of human cognition (Fodor & Pylyshyn, 1988): We can create an unlimited number of different ideas by freely combining content elements into new structures.

I assume that the DA region has a limited capacity that limits the complexity of structural representations that can be assembled by temporary bindings. The common variance reflected by measures of WM capacity (reviewed in Oberauer, 2005c) reflects essentially this capacity limit. The capacity limit arises from two sources, *retrieval competition* and *representational interference*. First, with an increasing number of context-to-content bindings, the attempt to retrieve one particular content element, cued by its context, suffers more competition from other contents bound to related contexts. The increasing degree of retrieval competition, however, cannot alone explain the steep decline of performance as the load on the DA region increases (Oberauer & Kliegl, 2001). Therefore, I assume a second form of interference by which the content representations bound in the DA region mutually degrade each other. Various candidate mechanisms for representational interference are currently explored through computational modeling, among them the superposition of content–context bindings in a common weight matrix (Farrell & Lewandowsky, 2002; Oberauer & Lewandowsky, 2008), and the overwriting of features shared by different contents in the DA region (Nairne, 1990; Oberauer & Kliegl, 2006; Oberauer & Lange, 2008).

2.1.3. The Focus of Attention

Reasoning not only involves the creation of structural representations but also their manipulation. For instance, we could use the array of three countries described above to play through what would happen if the population of China declined. We simply move the China token down along the dimension that represents population size while holding the other two elements in place—this might bring us to the conclusion that eventually the relative order of China and India will reverse while that of China with Russia stays the same. Selecting the target of a manipulation is the function of the focus of attention. It picks out one element among those currently held in the DA region and uses it as the object of a cognitive operation, in this case the operation of simulating a population decrease. My use of the concept of *attention* as applied to the contents of WM is based on an understanding of attention in purely functional terms, referring to a mechanism for the selection of representations for (cognitive) action (Allport, 1987).

This discussion of the focus of attention points back to a second function of the DA region. The contents of the DA region figure as the *selection set* for the focus of attention, that is, a small set of elements that are candidates for being brought into the focus. The focus of attention has direct access to these elements either through their content (e.g., when instructed to ''move India up'') or through their bindings to their places in the coordinate system (e.g., when asked ''which one is highest?''). Access to contents through their current contexts is what I mean by ''direct access.'' The role in a template, or place in a coordinate system, serves as a temporary address to which any arbitrary representation can be bound so that it can be accessed without any knowledge about its content.

2.1.4. Evidence for the Three Components of Declarative Working Memory

Consider the following task: You are asked to encode two short lists of digits for a short-term memory test, presented in two rows of boxes on a computer screen. After the digits are erased, one list is declared as temporarily irrelevant for the task. Next, you must perform a series of arithmetic operations on selected digits of the remaining list. The operations (e.g., ''+2'' or ''−4'') are displayed in one of the boxes of the relevant list, and you should retrieve the digit initially presented in that box, apply the arithmetic operation to it, and type the response as quickly as possible. This response is immediately followed by the next operation displayed in the same or a different box of the relevant list (see Figure 2). At the end of a series of operations, you are asked to recall both lists.

This paradigm (Oberauer, 2002) places clearly definable demands on all three components of declarative WM, and provides evidence for their
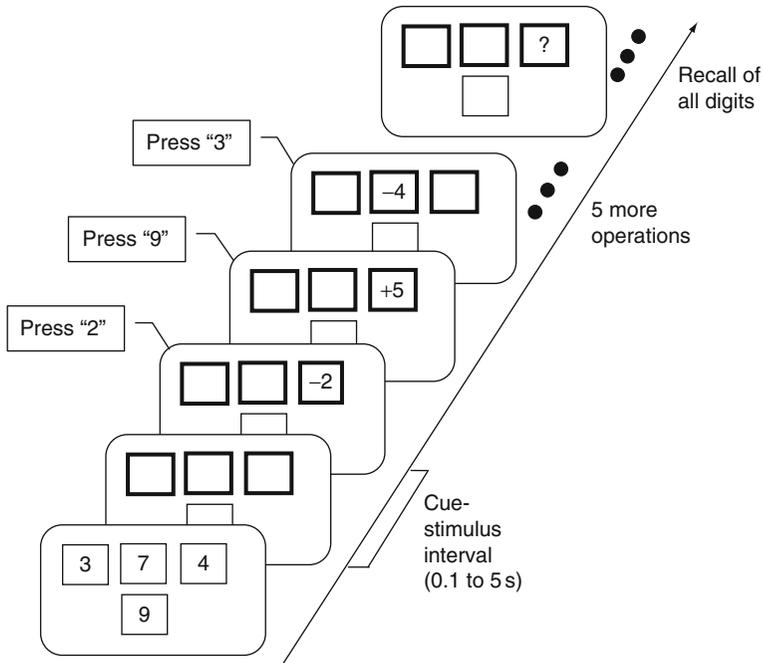
Figure 2   Sequence of events in a trial of Experiment 2 in Oberauer (2002). The trial has a relevant-set size of three and an irrelevant-set size of 1; the relevant set is cued by its frames turning red (illustrated here by thick lines). The third operation involves an object switch; the second is an object repetition.

distinction (for an illustration of how declarative WM applies to that paradigm, see Figure 3). Initially, all digits must be linked to their boxes by temporary bindings. These bindings make the digits directly accessible through their places in the cognitive coordinate system, that is, the boxes. Binding the digits to neighboring places in this episodic context integrates them into a structure in the DA region. The contexts of individual digits are not perfectly distinct, and therefore, cueing one of them—by placing an arithmetic operation in one box—also partially cues all other digits in the DA region, such that they all compete to some degree for being selected into the focus of attention. This competition leads to an increase of retrieval latencies with the set size of digits held in the DA region. This set-size effect is further compounded by representational interference (Oberauer & Kliegl, 2006). The results from my experiment were in line with these assumptions: When the first arithmetic operation was required immediately after the cue that designates the relevant list, response times to that operation increased with the length (or set size) of both lists (Oberauer, 2002).

Once people know which list will be used for the arithmetic operations, they can afford removing the other, temporarily irrelevant list from the DA
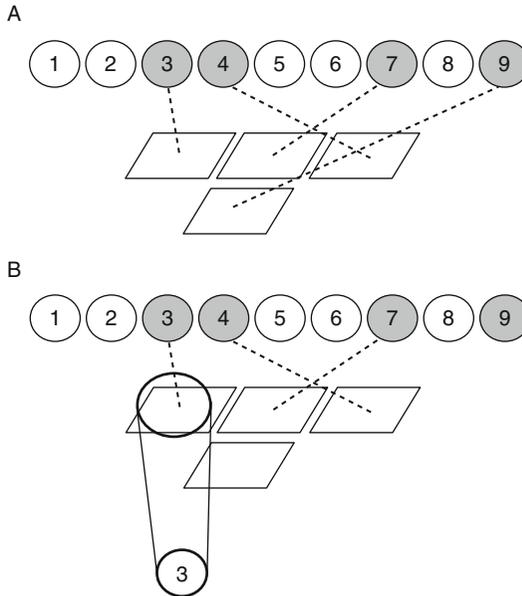
Figure 3 Declarative working memory operating on a trial of Experiment 2 in Oberauer (2002). Panel (A) illustrates the state immediately after encoding both lists. The small circles on top are representations of digits in LTM; those included in the present lists are activated (illustrated by shading). Digits are bound to their positions on the screen (thick dotted lines). Panel (B) reflects the state after the irrelevant list has been removed from the direct-access region—its bindings have been cut, but its content, the single digit 9, remains activated in LTM. The left-most digit of the relevant list must be accessed for processing, and the focus of attention (thick-lined oval) selects that location and uses it as a cue to retrieve the digit bound to it.

region, because they do not need to access its elements for the arithmetic task. Because that list will have to be recalled at the end of the trial, however, it must first be encoded into the activated part of LTM. This process takes about 2 s: Over the first 2 s after the cue that designates the relevant list, the set-size effect of the irrelevant list gradually drops to zero, while the set-size effect of the relevant list remains unchanged (Oberauer, 2002).

The WM system seems to have remarkable flexibility in swapping contents between the DA region and activated LTM. A list of digits or words that is not needed for processing only temporarily is still removed from the DA region within about 2 s, as reflected in the elimination of set-size effects. When that list is later cued as relevant for processing, it is brought back into the DA region—and the set-size effect on reaction times reappears (Oberauer, 2005b). With this list-switching procedure we can also measure the time it takes to switch between two list, that is, remove one from the DA region and replace it by another by retrieving it from activated LTM. In two experiments, I found these *list-switching costs* to

depend on the set size of the to-be-removed list, but not on the set size of the to-be-retrieved list (Oberauer, 2005b). The lack of a set-size effect for the list retrieved into the DA region suggests that this list is retrieved as a single chunk from activated LTM, an assumption to which I will return in the section on the interaction of WM with LTM.

Evidence for the focus of attention within the DA region comes from a further finding: When an arithmetic operation must be applied to the same digit as in the preceding step, latencies are shorter than when a different digit in the relevant set must be accessed. These so-called *object-switch costs* (Garavan, 1998; Oberauer, 2003) can be interpreted as the time it takes for the focus of attention to select a new element from the set held in the DA region. Supporting this interpretation, the object-switch costs increase with the length of the relevant list (Oberauer, 2003) but not with the length of the irrelevant list (Oberauer, 2002).

Object-switch costs are found in tasks in which an item from the memory set must be retrieved and used in a cognitive operation, and also in tasks in which the item is updated by replacing it with a new item (Oberauer, 2003). This shows that the objects that are selected into the focus of attention are not specific contents—when an item is updated, its content changes, but the object held in the focus remains the same, so that updating the same object again is faster than updating another object in the set. What remains constant across successive updates is the context by which the different items are distinguished and addressed (for instance, their boxes). Thus, what is selected by the focus is not the representation of a digit in semantic memory, but an episodic memory object defined by whatever is currently bound to a specific context, and this object maintains its identity across changes of its content. In this regard, objects in the region of direct access are similar to the *object files* in perceptual attention (Kahneman, Treisman, & Gibbs, 1992).

2.2. Procedural Working Memory

So far I have been concerned only with the representation and selection of the contents of cognitive activity, not the representation and selection of the cognitive operations themselves. I venture that the procedural part of WM can be conceptualized largely analogous to the declarative part, that is, as three embedded components that reflect three successive levels of selection of representations. A sketch of the architecture of WM, augmented by the procedural part, is given in Figure 4.

The representations in the procedural part are procedures. Procedures can be defined as (cognitive or motor) actions linked to the conditions in which they can be applied, and to the expected outcomes. Procedures form networks of associations in LTM, linking them sequentially and hierarchically to action plans (Schneider & Logan, 2007), and linking similar or
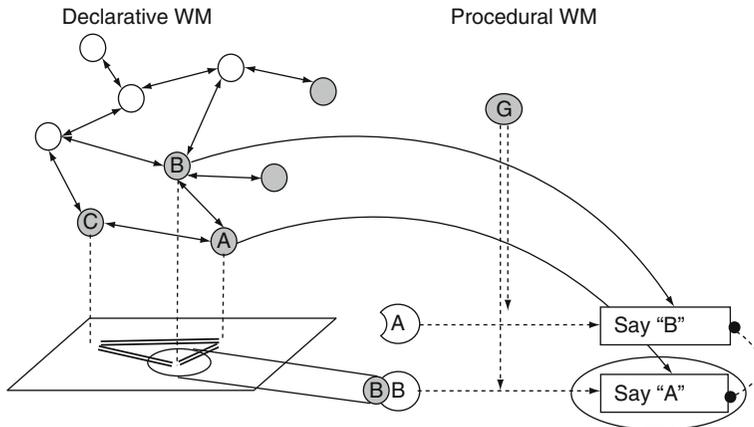
Figure 4   *Architecture of declarative and procedural working memory*. For explanation of the declarative side see Figure 1. On the procedural side, two condition-action bindings are illustrated that link representations of condition categories A and B to the response of saying ''B'' or ''A,'' respectively. These bindings implement the goal (G) to establish an incompatible stimulus–response mapping (i.e., saying ''B'' in response to A, and saying ''A'' in response to B). The two responses are connected by mutually inhibitory bindings (dotted line with nob ending). Together, these bindings constitute the current content of the bridge. The compatible mapping is represented by associations (bowed thin continuous lines) from declarative elements in LTM to the corresponding response representations; these associations are instances of procedural long-term memory. Executive processes, represented by thin broken lines, mediate between the goal G (which could be paraphrased as: ''respond to each letter by saying the other letter'') and the task set implementing it. In the present state of the system, the declarative element B is selected by the focus of attention. It matches the condition of one stimulus–response mapping in the bridge. The response bound to it is selected by the response focus (the second large oval).

equivalent procedures to each other. A subset of procedures is activated at any time; they form the *activated part of procedural LTM* (for evidence supporting the persistent activation of procedures in LTM see Woltz & Was, 2007). Procedures can be activated by representations of goals that they serve, as well as by stimuli to which they have been applied in the past (Waszak, Hommel, & Allport, 2003).

Activated procedural representations compete for control over a more central component of the procedural system, which I call the *bridge* (as in command bridge,[2] or as in bridge between stimulus and response). The bridge holds the currently operative task set, that is, the task set that is currently in control of thought and action. A task set is a coordinated set of

---

[2] One meaning of *bridge* is ''the platform on a ship from which the captain and officers direct operations'' (Oxford online dictionary, http://www.askoxford.com/).

condition–action–effect bindings (a.k.a. stimulus–response mappings) that specify for a confined set of conditions which action or cognitive operation to undertake, and which outcome to expect. The task set in the bridge is implemented as a set of directly executable procedures, such that whenever one of the conditions is represented in the focus of attention of declarative WM, the corresponding action is automatically executed. Thus, the operative task set acts as a "prepared reflex" (Hommel, 1998b). This procedural representation of a task differs from a declarative representation of the same task as a set of (typically verbal) instructions. Declarative representations of instructions can be instrumental in remembering what to do in which situation (Emerson & Miyake, 2003; but see Bryck & Mayr, 2005), especially early in practice, but I assume that immediate control of (cognitive) action lies with procedural representations, and a declarative instruction can be effective only by guiding the construction of a corresponding procedural representation in the bridge.

The bridge in procedural WM corresponds to the DA region in declarative WM, in that it serves to establish and hold temporary bindings between representations—here: between stimulus, response, and outcome representations. As in the DA region, these bindings are not limited to the learned associations in LTM. New bindings can be quickly and flexibly set up to link any stimulus category to any arbitrary response, for instance, to implement an experimental instruction. These bindings enable direct access from the stimulus to the response category (and its expected effect), as long as the particular set of bindings is upheld. The ability to form arbitrary *ad hoc* bindings between mental objects and cognitive operations on them is crucial for the flexibility of our thinking—we can make our responses to what we represent in declarative WM depend not only on the content of those representations but also on our current goals and intentions.

Finally, in the same way as the declarative WM system needs a mechanism for selecting a single object, the procedural system needs a mechanism for selecting a single (cognitive) action at any time. Corresponding to the focus of attention for objects, we must assume that the selected response attains a status that sets it temporarily apart from the competing responses in the task set. I will refer to this mechanism as the *response focus*.

The set of condition–action–effect bindings established as prepared reflexes in the bridge must be coordinated such that the conditions are mutually exclusive categories, each of which is uniquely mapped to a different action. This coordination is necessary to avoid cross talk from alternative mappings that could process the same objects as input but result in different, often incompatible operations (e.g., incrementing vs decrementing the population of China in the mental model discussed above). Therefore, the content of the bridge usually is a single task set that consists of a coordinated set of mutually exclusive condition categories, and the actions and effects bound to them. To optimize selection of a single response,

inhibitory links are set up between the response alternatives in the task set (Bogacz & Gurney, 2008; Usher & McClelland, 2001).

Tasks and procedures activated in LTM outside the bridge cannot directly control cognitive operations or actions—their responses are not candidates for selection by the response focus. They can, however, influence response selection in two ways. First, if the responses of an activated task set in LTM overlap with those of the task set currently established in the bridge, the former can contribute to *priming* these responses, that is, boost their activation without actually selecting them for execution (for a similar view see Lien & Proctor, 2002). Second, by virtue of being activated, a task set in LTM can be easily retrieved into the bridge (in particular when cued by either a task cue or by a stimulus that fits the condition part of the task set). In that case, the active task in LTM replaces the previous task set in the bridge and takes over control.

Both kinds of effects can be observed, for instance, in the Stroop paradigm: On incongruent trials, when the color word does not match the print color to be named, the irrelevant but strong association between printed words and speaking these words in LTM primes the response to speak the printed word. If the primed response is incongruent with the correct response (e.g., "blue" printed in green), correct responses are slowed because of the time needed to overcome the competition from the activated incorrect response. On some incongruent trials, however, people actually read the color word instead of naming the color. On these trials, the task to read printed words takes over the bridge, so that the unintended task is carried out—Kane and Engle (2003) refer to this as *goal neglect*. Evidence from a number of sources, reviewed by Kane and Engle, suggests that the two effects—slowed correct naming of incongruent colors and increased rate of word reading—are dissociable, consistent with my assumption that activated procedures in LTM can influence processing in two different ways, by priming responses and by displacing the current task set in the bridge.

2.2.1. Evidence for the Components of Procedural Working Memory

Evidence supporting the embedded component structure of procedural WM as outlined above is sparse, because procedural WM has rarely been addressed directly by empirical studies. Much work on perceptual attention and action selection, however, is relevant for procedural WM. In drawing on this literature, we must bear in mind one important difference: Whereas in studies of immediate memory, participants encode new content–context bindings with every trial, experiments on action selection typically provide an instruction mapping stimuli to responses at the beginning of the experiment and ask participants to use the same mapping for many trials. I believe that the latter experiments are nevertheless informative for procedural WM, because I assume that even well-practiced stimulus–response mappings must

be implemented as task sets in the bridge to control processing, in the same way as well–learned declarative associations in LTM must be projected into the DA region and maintained there by temporary bindings to be accessible for processing. Nevertheless, in typical experiments on action selection that use constant stimulus–response mappings it is difficult to rule out the theoretical alternative that, after a few initial practice trials, response selection is executed merely through gradually strengthened stimulus–response associations in LTM.

For this reason, a recent series of studies by Cohen-Kdoshay and Meiran (2007) is particularly important for bridging between the research traditions on WM and on action selection. These authors looked at the first few trials after instructing a new set of stimulus–response mappings and found that stimuli activated their corresponding responses automatically. This finding supports the idea that the task set is implemented as a set of ''prepared reflexes'' in WM, using *ad hoc* bindings between stimuli and responses, rather than emerging from gradual strengthening of associations in LTM over many trials. Thus, the concept of the bridge as a mechanism for *ad hoc* binding on the side of procedural WM, analogous to the DA region on the declarative side, gains some plausibility, although much more work is necessary to establish it firmly.

From the heuristic assumption that declarative and procedural WM work according to analogous principles we can make the prediction that the analogous components of both systems should show analogous empirical signatures. Table 2 lists the empirical signatures that I offered as evidence for the components of declarative WM above, and aligns them with corresponding phenomena in the domain of procedural WM. If this analogy holds, we should expect the following three empirical signatures in procedural WM: (1) Set-size effects should be observed for the number of stimulus–responses bindings held simultaneously in the bridge, but not for stimulus–response mappings in task sets merely activated in LTM. (2) Switching between task sets takes time, and these switch costs increase with the set size of the task set switched away from, but not with the set size of the new task set. (3) Repeating a response selected previously should be faster than switching to a new response. A fourth prediction, which so far has no direct analogue on the declarative side, can be derived from the assumption that the response focus holds only one response representation at a time: Procedural WM can select only one response at a time, and this limitation creates a bottleneck for cognitive operations. I now discuss evidence for these four predictions (the evidence for the first two is closely linked, so I discuss it jointly in the following section).

2.2.1.1. Set-Size Effects for Stimulus–Response Mappings and Task-Set Switching Costs
Relevant evidence for the first two predictions comes from the task switching literature (for a review see Monsell, 2003).

Table 2 Analogous Phenomena in Declarative and Procedural Working Memory.

Phenomenon	Declarative WM	Procedural WM
Set-size effects	Set-size effect for relevant set of objects in DA region No set-size effect for currently irrelevant set in activated LTM (Oberauer, 2002)	Set-size effect for number of S–R bindings in operative task set in the bridge No set-size effect for currently not operative task set (Hübner et al., 2004)
Switching structural representations	Costs for switching between lists of items Lists-switch costs increase with set size of old list but not with set size of new list (Oberauer, 2005b)	Costs for switching between task sets Switch costs increase with number of S–R bindings in old task set but not with number of S–R bindings in new task sets (partial support in Hübner et al., 2004)
Switching focus	Switch costs for focusing on new object, increasing with set size in DA region (Oberauer, 2002, 2003)	Response switch costs (Bertelson, 1965)
Serial selection	Only one object can be selected into the focus at the same time	Only one response can be selected at the same time (Pashler, 1994), but there are exceptions (Oberauer & Kliegl, 2004)

When people switch between two tasks, only the currently relevant task set is held in the bridge, other task sets that were recently used and are expected to be used in the near future will be activated in procedural LTM. Switching from one task set to another involves removing the old task set from the bridge, and retrieving a new task set from LTM (Mayr & Kliegl, 2000). Task-switch costs reflect to a large degree the time for this process.[3]

One important distinction between activated LTM and the DA region in declarative WM is that only representations in the DA region compete for selection into the focus, and therefore, set-size effects on access latencies are found only for the set in the DA region, not for any additional sets held in activated LTM. By analogy, we should expect a set-size effect for the number of stimulus–response mappings in the currently selected task set, but not for other active but not selected task sets. There is no study directly addressing this prediction, but some relevant evidence can be found in a couple of experiments by Hübner, Kluwe, Luna-Rodriguez, and Peters (2004). In their Experiment 1A, people switched between a task set with two stimulus–response mappings and a task set with four mappings. Unsurprisingly, reaction times were about 200 ms longer in the task with four mappings. In Experiment 1B, people switched between two tasks, both of which had four mappings. Reaction times were about as large as in the four-mapping task of Experiment 1A (although no formal comparison was made between experiments). Thus, the number of mappings of the relevant task set—which must be established in the bridge—had a large influence on reaction times, but the number of mappings in the currently irrelevant task set—which arguably is held activated in LTM, but outside the bridge—has little or no effect on reaction times.

Task-switching situations typically engender a general cost from holding more than one task set ready for use, compared to a situation where only a single task set is relevant for a whole block of trials. These "global switch costs" or "mixing costs" could be thought of as evidence against the contention that only one task set is held in the bridge at any time in task-switching paradigms. Rubin and Meiran (2005), however, have shown that mixing costs do not reflect an increased load on a limited-capacity WM. If all task sets involved in a switching experiment added to the load on the bridge, increasing the number of task sets should increase mixing costs. This was not the case in Rubin and Meiran's study (for further evidence see Kray, Li, & Lindenberger, 2002). Mixing costs increased, however, when the stimuli were "bivalent," that is, not only the currently relevant but also the currently irrelevant task set(s) could be applied to them. This suggests that the currently irrelevant task sets are maintained outside the bridge, but in a

---

[3] Task-switching costs have been decomposed into several components (Meiran, Chorev, & Sapir, 2000). I use the term task-switch costs to refer to the difference between switch trials and no-switch trials on mixed blocks at short cue-stimulus intervals (i.e., "switching costs" in the terminology of Meiran et al., 2000, p. 248).

highly activated state that makes them easy to retrieve back into the bridge, in particular when they are cued. Other experiments have shown that stimuli fitting the conditions of the condition–action links in a task set act as retrieval cues for these task sets when they have been used recently (Waszak et al., 2003). Mixing costs could thus be interpreted as arising from occasional inadvertent retrieval of a currently irrelevant task set into the bridge, or from the extra time it takes to prevent such retrieval. When the responses of the task sets overlap, priming of the wrong responses along stimulus–response associations outside the bridge can additionally contribute to mixing costs—this factor seems to play a significant role particularly in older adults (Mayr, 2001).

Further evidence for the assumption that only relevant responses contribute to response competition comes from experiments with the Stroop paradigm (for a review see MacLeod, 1991). Color words interfere with naming the word's ink color only if the word refers to a color that can also occur as an ink color. This finding can be explained by assuming that the bridge holds stimulus–response mappings only for the colors that are expected to occur as ink colors. Color words not occurring as ink colors therefore cannot prime a response that belongs to the selection set in the bridge. Roelofs (2001) reviews results with a variant of the Stroop paradigm, picture naming in the context of interfering words. When the set of pictures used is small and frequently repeated, words not in that set do not interfere. With a large and rarely repeated set of pictures, however, all words interfere with picture naming. Roelofs explains this observation by assuming that a selection set of responses can be maintained in short-term memory only when it is sufficiently small and frequently repeated. In the present framework, this ''short-term memory'' would be the bridge of procedural WM. But why does any word interfere with picture naming when the set of pictures is large? In such a situation, the task set in the bridge cannot specify all stimulus–response mappings individually. Rather, the task set must include a call to declarative LTM. Expressed as a verbal rule, the task set would be ''If a picture is presented, retrieve its name from LTM and say it aloud.'' Interference from the distracting word arises because the word activates a corresponding representation in declarative LTM, which interferes with the retrieval of the name for the picture.

I assume that set–size effects in the central components of WM—the DA region and the bridge—arise not only from competition for selection, but also from representational interference. The latter causes increasing degradation of representations with larger set size, such that only a limited number of elements can effectively be held in the central components at any time. Whereas there is plenty of evidence for such a capacity limit for declarative WM, there is very little on the procedural side. Some initial support for a capacity limit on procedural WM comes from a study by Ellenbogen and Meiran (2008). They created a dual-task situation in which

two speeded choice tasks had to be performed in rapid succession, and the response of the second task (saying a color word) could be compatible or incompatible with the stimulus of the first task (classifying the stimulus by its color). In dual-task setups like this, *backward compatibility effects* are frequently observed (Hommel, 1998a), showing that the response to the second task is at least activated before the response to the first task is selected. Ellenbogen and Meiran hypothesized that in this case, the cognitive system tries to hold both task sets simultaneously in a capacity-limited WM. Doing so would be more difficult, however, when the load on WM by the first task is raised by increasing the number of stimulus–response mappings. Ellenbogen and Meiran therefore predicted that the backward-compatibility effect should disappear, and this is what they found when they increased the number of mappings in the first task to six.

The finding of Ellenbogen and Meiran (2008) can be contrasted with the finding of Kiesel, Wendt, and Peters (2007) in a task-switching paradigm in which both tasks are mapped to the same set of response alternatives. In this paradigm, congruent stimuli that are mapped to the same response by both tasks are typically responded to faster than incongruent stimuli that are mapped to different responses by the two tasks. Kiesel et al. found that this *congruency effect* was not diminished by increasing a concurrent load on declarative WM. The congruency effect was modulated, however, by the relative frequency of stimulus–response pairings in previous trials. These results suggest that, different from dual-task paradigms, task-switch paradigms do not encourage people to hold two task sets in the bridge at the same time; rather, congruency effects in task switching reflect the strengthening of stimulus–response associations of both tasks in LTM, such that each stimulus primes the responses for both tasks. Further evidence that the congruency effect arises from task sets that are activated in LTM has been obtained by Meiran and Kessler (2008). They showed that the congruency effect occurs only if the task set creating involves an abstract representation of the response categories (e.g., "up," "down") that is well established in LTM either preexperimentally or through training within the experiment. When no such representation exists in LTM, none can be activated, and without mediation through an activated abstract response code the stimulus representation cannot prime the corresponding response.

The second empirical signature derived from the analogy with declarative WM includes the prediction that task-switch costs depend on the set size of the task set switched away from, but not on the set size of the task set switched to. Some evidence supporting this prediction can again be found in the experiments by Hübner et al. (2004). They asked people to switch between a task with few stimulus–response mappings and a task with many stimulus–response mappings and found that task-switch costs were smaller when switching to the task with the larger number of stimulus–response mappings. In these experiments, switching to a task set with more mappings

implied switching away from a task with fewer mappings, so these findings are at least consistent with the hypothesis that switch costs increase with the set size of the task switched away from, and not with the set size of the task switched to. A direct test of this hypothesis would require fully crossing the set sizes (i.e., number of mappings) of both tasks, analogous to my experiments on list-switch costs (Oberauer, 2005b).

In the context of set-size effects in procedural WM, and the capacity-limiting factors they reflect, the work of Duncan and colleagues on *goal neglect* is also relevant (Duncan, Emslie, Williams, Johnson, & Freer, 1996; Duncan et al., 2008). In their paradigm, participants see a series of pairs of stimuli (e.g., two letters), each screen displaying one stimulus on the left and one on the right side. Participants must respond to the stimuli on one side only (e.g., name the letter on the left side). At some point in the series, a symbol is presented centrally instead of a pair of letters, and this signals a side change; from now on the letters on the right side must be named. The whole series is presented at a fast pace. Goal neglect is demonstrated when people ignore the side-change stimulus, thus continuing to respond to the stimuli on the same side as before. Goal neglect occurs even when people can recall at the end of the experiment the instructions on how to respond to the side-change stimulus. This finding implies that goal neglect does not simply arise from forgetting of the relevant instruction. Further, general intelligence is negatively correlated with the frequency of goal neglect.

Duncan et al. (2008) showed that the probability of goal neglect depends not on the complexity of the main task actually conducted, but on the complexity of the task that people are instructed and given practice on. For instance, in their Experiment 3 one group of participants was instructed, and given brief practice with, a task-switching protocol by which they alternated between reading letters and adding digits on the relevant side of each screen, whereas another group was initially instructed on only one of these tasks. After the practice trials both groups were told that they were going to work on only one task (e.g., only letter reading). The group initially instructed on the task-switching protocol committed more goal neglect. Duncan et al. interpret their results as evidence for a capacity limit on the ability to establish a *task model* based on instructions. When the task model's complexity exceeds capacity, it is simplified by dropping some part of it, for instance, the side-change rule.

In the framework I propose here, Duncan's goal-neglect paradigm can be understood as a task-switching paradigm. Work on each trial starts with the task set for the main task (e.g., letter reading) in the bridge. When the side-change symbol comes up, people must rapidly switch tasks, that is, retrieve the side-change task set from procedural LTM into the bridge; execution of that task set shifts visual attention to the other side. On this analysis, goal neglect can be understood as a failure to retrieve the task set for the side change. Given the high time pressure in this paradigm, such failure

becomes more likely in people who process the main task more slowly, because they will still be busy completing the last operation of the main task when the side-change symbol appears, and because general intelligence is correlated with processing speed, this can explain the correlation of goal neglect with intelligence. Successful retrieval of the side-change task set also depends on an association between the retrieval cue (i.e., the side-change symbol) and the task set, which can be established during the practice trials, or through mental practice in response to the instructions. The effect of instruction complexity could arise because people who received more complex instructions suffer from more representational interference in the bridge during practice and therefore are less likely to accomplish the practice trials correctly. In particular, they are less likely to do the side-change correctly in the practice phase. As a consequence, they have weaker associations between the side-change symbol and the side-change task set, and thus are more likely to commit goal neglect.

2.2.1.2. Response-Switch Costs   It has long been known that repeating the same response in a series of choice reaction-time tasks results in a benefit of about 100 ms, even when the stimulus triggering that response changes between trials (Bertelson, 1965). This response-repetition benefit, or response-switch cost, can be regarded as the analogue to object-switch costs, or object-repetition benefits, in declarative WM. Response-repetition benefits seem to be smaller than object-repetition benefits, and they do not last long, suggesting that the response focus does not hold on to a response representation after it has been executed as long as the focus of attention holds on to a declarative representation after it has been used in a cognitive operation.

Recent research on the response-repetition benefit has revealed that it disappears, or is even reversed, when the task set used to select the response, or the stimulus category demanding the response, is changed (Kleinsorge, 1999; Schuch & Koch, 2004). These findings imply that the representation held in the response focus cannot be a representation of the physical response (e.g., a motor program or an action plan) but rather must be a representation of the "meaning" of the response (Schuch & Koch, 2004). For instance, when an odd–even judgment on digits is followed by a large–small judgment, pressing the left key could mean "odd" on the first trial, but mean "large" on the second trial. Therefore, repeating the same physical response does not incur a repetition benefit. I conclude from these findings that the representation in the response focus is a representation of the decision or conclusion that the procedural system has arrived at by processing a declarative representation through the currently operative task set. That decision or conclusion can be, but does not have to be, linked to a physical action. A physical action resulting from a decision, and its expected outcome, must be part of the task set that controls overt actions, and there

must be a mechanism to select the appropriate physical action over competing alternatives, but that mechanism apparently does not hold on to the selected representation after its execution, and therefore does not give rise to repetition benefits for the actions themselves when their meaning changes.

2.2.1.3. Response-Selection Bottleneck The fourth prediction motivated by the analogy between declarative and procedural WM is that the response focus selects only one response at a time. Evidence for this prediction comes from the dual-task literature, which strongly points to the existence of a response-selection bottleneck for even very simple choice tasks (Byrne & Anderson, 2001; Pashler, 1994). Evidence for a bottleneck also exists for purely cognitive operations that update representations in WM (Oberauer & Göthe, 2006). The evidence that objects in declarative WM can be updated only one at a time suggests that not only responses in the response focus but also objects in the focus of attention are selected one at a time, thus providing preliminary evidence for a bottleneck on object selection that is analogous to the bottleneck on response selection (see the bottom-left cell of Table 2).

Neither the presumed object-selection bottleneck nor the response-selection bottleneck, however, is immutable. Using a memory-updating paradigm similar to that of Oberauer and Göthe (2006) but with reduced memory load, my colleagues and I found that young adults (but not old adults) can, with substantial practice, acquire the ability to update one digit and one spatial position in WM simultaneously without mutual interference between these processes (Göthe, Oberauer, & Kliegl, 2007; Oberauer & Kliegl, 2004).

The theoretical framework developed here provides some guidelines for explaining why, and under which circumstances, overcoming the selection bottlenecks should be possible. The bottlenecks are assumed to arise not from hard-wired architectural constraints but from the functional constraints on the attentional mechanisms in WM. The function of these mechanisms is to unambiguously select an object for processing in declarative WM, and a response in procedural WM. Limiting the contents of the two foci to a single representation is a straightforward way of avoiding selection ambiguity, and likewise, limiting the content of the bridge to a single task set is an effective way of avoiding cross talk between task sets (cf. Logan & Gordon, 2001). These constraints, however, can be relaxed in situations where ambiguity and cross talk are highly unlikely. A digit and a spatial position are unlikely to be confused. Moreover, in our experiment people only held one digit and one spatial position in declarative WM at any time, thus minimizing the selection demand on the focus of attention because of a lack of competitors. Under these conditions, the focus of attention can hold a digit and a spatial position jointly without creating

selection ambiguity. Likewise, an arithmetic computation and a shift of an object in space are unlikely to create cross talk because the shift cannot be applied to the numerical value of the digit, and the spatial location cannot be affected by an arithmetic operation. Therefore, the system can, with sufficient practice on the specific task combination, learn to relax its default cautionary setting and hold two objects in the focus of attention and two task sets in the bridge, and select two responses into the response focus simultaneously.

According to this view, whether parallel processing is possible depends not on the individual tasks but on the distinctiveness of representations and processes in a task combination. Consistent with this view, practice on the task combination, not practice on the individual tasks, enabled parallel processing (Oberauer & Kliegl, 2004). Moreover, in an unpublished follow-up study combining the digit updating task with a letter updating task, people could not perform both tasks in parallel even after extended practice (Brambosch, 2003). This finding shows that the distinctiveness of the two task sets is an important prerequisite for acquiring parallel processing.

2.2.2. Executive Processes

Speaking about WM one cannot avoid speaking of executive processes, because the two terms have been so closely linked in the literature. Some use the terms *executive processes* or *executive functions* in a very encompassing way that treats the whole of WM as one instance of them. Nothing is to be gained for an understanding of WM from using that concept. Others use the term *central executive* to refer to the processing mechanism(s) of WM, setting them apart from the storage mechanisms (e.g., Goldman-Rakic, 1997). Used in this way, the term roughly coincides with my concept of procedural WM, and that is why I placed my discussion of executive function in this section. It is worth noting, however, that Baddeley (1986) originally modeled his central executive after the Supervisor Attentional System (Norman & Shallice, 1980), which is meant to supervise and control ongoing cognitive processes rather than execute them itself.

In keeping with the general idea of Norman and Shallice, I find it useful to distinguish between primary processes and executive processes. I call *primary processes* those that take declarative representations (either from perception or from declarative WM) as input and produce manipulations of declarative representations or overt actions as output. In contrast, I call *executive processes* those that have as their output a manipulation of the conditions of primary processes. Executive processes can change the conditions of primary processes by changing their parameters (Logan & Gordon, 2001), such as changing the speed-accuracy trade-off criterion. Alternatively, executive processes can control primary processes by changing the representations in declarative or procedural WM that guide the primary processes, such as manipulating the activation levels of

representations in LTM, and updating the contents of the DA region or the bridge.

Individual–differences research has distinguished three categories of executive function tasks: *updating of WM, inhibition of prepotent responses*, and *task-set switching* (Friedman et al., 2006; Miyake et al., 2000). All three categories of tasks involve establishing a new representation that is necessary for an intended task against a conflicting old representation that is irrelevant or misguiding. Whereas *updating* refers to replacing old representations in declarative WM by new ones, *inhibition* and *switching* refer to establishing new representations in procedural WM in competition with old ones. The difference between inhibition and switching is that inhibition tasks require establishing a new task set in the bridge against a habitual response tendency intruding from procedural LTM, whereas switching requires establishing a new task set against another task set that is currently held in the bridge, but does not necessarily have a strong representation in LTM. Thus, inhibition requires implementing the intended task set in the bridge strongly, so that a highly activated but unintended procedure in LTM is not inadvertently retrieved into the bridge and takes over control. Task–switching involves replacing the current task set in the bridge by a new one; thus, it can be conceptualized as updating procedural WM.

Table 3 presents a proposed taxonomy of demands on declarative and procedural WM, as reflected in tasks used to study WM or executive functions. The common denominator of these tasks is that they require maintaining goal–relevant representations in a state of immediate avail-ability. The representations whose maintenance is performance limiting can be either declarative (in the first column) or procedural (in the second column). The first row includes tasks that are primarily limited by the ability to maintain representations available, without a large role for competition. The tasks in the second row, *Overcoming Competition*, add to this the requirement of overcoming distraction by competing representations that become strongly activated by misleading stimuli, misleading associations, or both. The third row, *Updating*,[4] combines maintenance with the require-ment to update representations. This class of tasks is interesting because they focus on a fundamental dilemma of WM, to strike the balance between holding on to representations and replacing them by new ones (Dreisbach & Goschke, 2004; Durstewitz & Seamans, 2002).

The tasks regarded in the literature as reflecting executive functions are found in the second and third row of the table, reflecting *Overcoming Competition* and *Updating*, respectively, in addition to the maintenance demand common to all three rows. The demands in the bottom two rows

---

[4] I use the term *Updating* in capital letters to refer to the category of tasks in the bottom row of Table 3, which involve updating of either declarative or procedural WM, and *updating* in lower-case to refer to the updating factor in Miyake et al. (2000), which includes only updating of declarative WM.

Table 3 Taxonomy of Demands on WM.

Demands	Declarative WM	Procedural WM	Common required function
Maintenance	Short-term recall or recognition *Digit span, Sternberg recognition paradigm*	Hold new task set to guide action *Speeded choice RT with arbitrary S–R mapping*	Activation, binding
Maintenance + Overcoming competition	Overcoming proactive interference from LTM *PI paradigm, recent probes task,[a] modified Sternberg task,[b] storage + processing paradigms[c] (?)*	Overcoming prepotent responses ("inhibition") *Stroop task, antisaccade task*	Strong binding, executive processes
Maintenance + Updating	Replacing old contents of declarative WM by new ones ("updating") *Keep-track,[d] memory-updating[e]*	Replacing old task sets by new ones ("switching") *Task-set switching paradigm*	Flexible binding and unbinding, executive processes

Note: Example tasks are given in italics; brief descriptions of not commonly known tasks follow.

[a] The recent-probes task is a variant of the Sternberg recognition task in which negative probes of the current trial were included in memory lists on recent trials, thus producing misleading familiarity.

[b] In the modified Sternberg task, participants encode two lists, one of which is post-cued to be forgotten. Recognition probes from the to-be-forgotten list create misleading familiarity.

[c] Storage + processing paradigms (e.g., reading span, operation span) combine immediate recall of short lists with episodes of a typically unrelated processing task either in between or after presentation of list items. The status of storage + processing paradigms is uncertain: If the concurrent processing task leaves traces in WM that compete for recall with the to-be-remembered list items, the task reflects overcoming competition; otherwise it just reflects maintenance.

[d] In keep-track, participants see a list of nouns and must remember the last noun in each of a number of categories.

[e] In memory updating, participants remember several digits and update them individually by arithmetic operations.

have in common that they require the frequent intervention of processes that take goal representations as input and generate strong goal–appropriate representations in WM as output, in other words, executive processes. There is, however, an important difference between *Overcoming Competition* and *Updating*: Performance in tasks that require overcoming competition must be assumed to depend on two factors, the effectiveness of executive processes in establishing the goal-relevant representation, and the ability of the WM system to hold and maintain strong bindings that can overcome prepotent memories or response tendencies from LTM. The ability to maintain strong bindings is, in my view, the main limiting factor of WM capacity. Thus, tasks measuring the ability to overcome competition reflect a mixture of two variables, the efficiency of executive functions and the capacity of WM. In Updating tasks, in contrast, having high WM capacity is unlikely to be helpful. A person with high capacity would have the disadvantage of having to overcome strong bindings to remove the old content of the DA region or the bridge, but the advantage of being able to strongly bind the new content. A person with low capacity would have to overcome weaker bindings with a weaker mechanism to establish new bindings. There is no reason to assume that one or the other will be more effective in updating. In other words, whereas in overcoming-competition tasks, the WM system fights conflicting tendencies from external sources, in updating tasks, WM fights itself. Therefore, Updating tasks could be regarded as a relatively pure reflection of executive functions because performance is unlikely to be affected by WM capacity. The only determinant of success in Updating tasks should be the speed and accuracy of executive processes that decide on and carry out the updating.

Do executive processes have a special status in the cognitive system? One might argue that they do not. In this view, executive processes use the same mechanisms as primary processes, they just happen to have responses that control the cognitive system itself rather than some aspect of the environment. A consequence of this view is that executive processes must share with primary processes the limited capacities of declarative and procedural WM. In particular, the bottleneck in the bridge would force the system to alternate between task-sets for primary processes and task-sets for executive processes. This idea raises the question how the switch between primary and executive task sets is managed—it seems that an executive process is needed to switch from a primary to an executive process. This looks like very bad design.

For an efficient WM system we need to assume that at least some executive functions avail of a separate mechanism that enables them to run in parallel with primary processes, so that they can intervene in primary processes. This assumption is in agreement with theory and data on at least one executive process, stopping an action. Logan and Cowan (1984) analyzed the stop-signal paradigm, in which participants do a speeded choice

task and are instructed to stop responding when they perceive a stop signal, which occurs on a minority of trials. Logan and Cowan successfully applied a race model to data from this paradigm in which the primary response-selection process and the stop process run in parallel.

Not all executive processes, however, seem to have the privileged status of being independent of primary processes. Monitoring of one's performance, and correction of errors, can be regarded as a prime example of an executive function. It has long been known that when people commit an error in a series of simple choice tasks, they slow down on the following trials (Laming, 1968). This posterror slowing has been interpreted as reflecting an adjustment of the speed-accuracy trade-off setting to avoid further errors. A recent study by Jentzsch and Dudschig (2009), however, casts doubt on this interpretation. They showed that posterror slowing is more likely a manifestation of a processing bottleneck: After an error, processes interpreting the error and making adjustments occupy the processing bottleneck for some time, thus postponing processing of the next stimulus. It seems that at least some executive processes compete with primary processes for the bridge in procedural WM.

To conclude, an analysis of functional necessities, as well as some empirical evidence, point to the existence of a separate mechanism for executive processes in addition to declarative and procedural WM, such that executive processes can run in parallel with primary processes. This mechanism, however, seems not to be used by all processes that are regarded as executive in the literature, and that match my definition proposed above. Future research will hopefully distinguish between executive processes that must share mechanisms with primary processes and others that do not, and this might provide a more robust basis for classifying cognitive processes than the crude distinction into primary and executive processes that I can offer at this point.

2.2.3. Are Declarative and Procedural Working Memory Separate Systems?

I have assumed so far that, at least for the central components of the WM system, declarative and procedural WM are parallel but separate systems. An alternative is that there is only one WM system that is used for both declarative and procedural representations. Thus, there would be no significant difference between bindings linking memory objects to their contexts, and bindings linking objects to (cognitive) actions to be performed on them, and the two kinds of bindings would share the same limited capacity, that is, they would interfere with each other. Moreover, there would be a single focus of attention that selects a composite of a mental object and an associated cognitive operation to be performed on it. Thus, objects and operations would not be selected independently but as tightly packed units. This latter view is suggested by the Theory of Event Coding (Hommel,

Müsseler, Aschersleben, & Prinz, 2001), which proposes that perceived objects and events and the actions we perform (or plan to perform) on them are represented as integrated *event codes*.

One straightforward prediction of the independence assumption is that increasing the load on declarative WM should not affect the efficiency of executing a task set held in procedural WM, and conversely, increasing the load on procedural WM should not impair retention of information in declarative WM. This prediction is complicated somewhat by the fact that, whereas task sets are units of procedural knowledge, task *instructions* are instances of declarative knowledge. Often, in particular for novel and not yet practiced tasks, people hold the (usually verbal) task instruction in declarative WM to back up the task set that implements that instruction. Increasing the load on declarative WM could disrupt memory for the task instruction, and as long as instruction memory contributes to task performance, the latter would be impaired. Likewise, making the task set more complex usually implies making the instruction more complex, and thus a more complex task would create more interference with other contents of declarative WM if the person attempts to hold the task instruction in declarative WM. Therefore, an informative test of the above prediction would have to involve a well-practiced task for the manipulation of procedural complexity, which no longer relies on declarative instructions. Alternatively, a nonverbal WM load that does not interfere with verbal instructions held in declarative WM could be used to manipulate declarative WM load. As far as I am aware, no such study has yet been conducted.

Results from individual-differences research turn out to be more informative. Schmiedek, Oberauer, Wilhelm, Süß, and Wittmann (2007) analyzed the relationship between measures of (declarative) WM capacity and reaction times in speeded two-choice tasks with arbitrary stimulus–response mappings. The choice tasks arguably rely at least in part on the robustness of stimulus–response bindings in procedural WM. Schmiedek et al. applied the diffusion model (Ratcliff, 1978) to the reaction time distributions to isolate the decision component from other components of variance in reaction times; the decision component presumably reflects the strength of stimulus–response bindings. In a structural equation model, a latent factor for the decision component correlated highly, but not perfectly, with a latent factor reflecting WM capacity. This finding suggests that the ability to maintain bindings in declarative and in procedural WM share a substantial amount of variance. Further support for this conclusion comes from a study by Wilhelm and Oberauer (2006), who related speed in four-choice tasks with compatible and with arbitrary stimulus–response mappings to measures of (declarative) WM capacity. Choice tasks with arbitrary mappings yielded much larger reaction times than those with compatible mappings. A factor reflecting individual differences in the size of the compatibility effect was

highly correlated with WM capacity. This observation is again consistent with the view that there is a common source of variance in the ability to maintain temporary bindings in declarative and in procedural WM. It should be added, however, that a common source of variance does not imply a common mechanism. It is possible that declarative and procedural WM are separate mechanisms that are affected by the same source of individual differences (e.g., global neural noise affecting different neural networks in the same way).

Evidence pointing more toward a distinction between declarative and procedural WM comes from a series of studies of Miyake and colleagues (Friedman et al., 2006, 2008; Miyake et al., 2000). In these studies, the *updating* factor and the *task-set switching* factor were always distinct, though correlated. According to the taxonomy in Table 3, these two factors reflect the *Updating* demand for declarative and procedural WM, respectively. Further, Friedman and Miyake (2004) found that a latent factor reflecting resistance to proactive interference was uncorrelated with their *inhibition* factor. These two factors could be interpreted as representing the *Overcoming-Conflict* demand in declarative and procedural WM, respectively, and their lack of correlation would reflect a dissociation between declarative and procedural WM. The latter result is ambiguous, however, because the proactive-interference measure was obtained from delayed recall tasks that arguably reflect more the robustness of new LTM associations than the robustness of temporary bindings in declarative WM.

A brain-imaging study by Nelson, Reuter-Lorenz, Sylvester, Jonides, and Smith (2003) provides more direct evidence for a dissociation of mechanisms for overcoming conflict in declarative and procedural WM. Using the Sternberg recognition paradigm, they created conflict in declarative WM by presenting probes that were not in the current memory list but in the preceding memory list. These recent negative probes create a conflict between familiarity and recollection. Overcoming the misleading familiarity signal requires strong recollection, which rests on strong bindings between memory items and their list contexts (cf. Oberauer, 2005a). In addition, Nelson and colleagues created conflict in procedural WM by presenting negative probes that were positive probes on the preceding trial, such that the correct response on the present trial conflicts with the correct response on the preceding trial to the same stimulus. Overcoming response conflict arguably requires a strong task set to minimize the influence of the stimulus–response association carrying over from the preceding trial. Nelson et al. (2003) found that conflict between familiarity and recollection increased activation in prefrontal cortex, whereas response conflict increased activation in anterior cingulated cortex. This result suggests that different brain circuits are involved in resolving conflict in declarative and procedural WM.

Overall, there is surprisingly little evidence speaking directly to the degree of independence of declarative and procedural WM. My preference

for the assumption of two separate systems, though consistent with extant findings, is motivated mainly on theoretical grounds: Being able to select declarative representations (i.e., the objects of thought) and procedural representations (i.e., our actions on them) independently seems a desirable feature for a WM system, because it enables us to select an object of thought without deciding at the same time on how to mentally manipulate it or how to act on it. Having separate representational systems and selection mechanisms for declarative and procedural representations is a straightforward way of realizing this independence.

3. Working Memory and Long-Term Memory

There is an increasing recognition that the processes attributed to WM are rarely free from contributions from LTM (Burgess & Hitch, 2005; Unsworth & Engle, 2007). In my view, there are two ways in which LTM contributes to the functioning of WM, through activation of existing representations, and through learning of new structural information that can be retrieved back into WM.

3.1. Activated Long-Term Memory

I see three roles of activated representations in LTM for processes that are usually attributed to WM, such as immediate recall or recognition. The first role is to increase the efficiency of processing new stimuli matching already activated representations. This role is manifest in short-term priming (McKone, 1998). In addition, the perceived increase in processing fluency for stimuli matching already activated representations seems to be an important ingredient of their familiarity (Whittlesea et al., 1990).

The familiarity signal generated in response to a stimulus can be used as one source of information for recognition decisions because familiarity is correlated with recency of use. For short-term recognition (i.e., recognition decisions following within seconds after initial encoding of the memory items), processing fluency, and thus familiarity, will be mainly determined by the degree of activation of the memory items in LTM. A further source of information for recognition is recollection, that is, retrieval of episodes of encountering the stimulus before in a relevant context (e.g., the context of the last seen memory list). For short-term recognition, recollection implies retrieving memory items from the DA region, which provides bindings between the items and their contexts. The distinction between activated LTM and the region of direct access receives support by studies of short-term recognition that fit this brief sketch of a dual–process model of recognition (Oberauer, 2001, 2008; Oberauer & Lange, 2009).

A second role for activation in LTM is that activated representations narrow down the search set for reconstruction of memory traces in the DA region. Representations in the DA region can be degraded by representational interference, and in that case retrieval involves recovery of the complete representation by retrieving the best-matching representation in LTM—this form of pattern completion is often referred to as *redintegration* (Lewandowsky, 1999). Having a limited set of activated candidate representations makes redintegration faster and less error prone. Using the ordering of countries by population as a simplified example: Imagine that the representations of countries in the DA region consisted only of their names, and due to interference, at some point the structure consists of degraded memory traces reading: "I_D__, CH__A, R_S___." Matching these traces against the whole vocabulary of the language would make recovery very difficult, but having a search set limited to countries renders the task easier.

A third role for activation in LTM is to provide a limited mechanism for representing serial order. A list of items can be represented in order by imposing a gradient of activation on them, declining from the first to the last. The items can be reproduced in that order by retrieving the one with the highest activation and then suppressing it, thus leaving the next in line with the highest activation. This mechanism has been used in models of forward serial recall of lists (Grossberg & Stone, 1986; Page & Norris, 1998).

Activation gradients, however, are limited in that they do not provide direct access to a specific element—to access the fourth element in a list, for example, retrieval has to proceed from the beginning until it reaches the fourth element. Activation gradients also cannot represent multidimensional structures, such as the location of various elements in two-dimensional space. A powerful mechanism for representing structure requires the ability to establish links between representations. Temporary bindings establish these links in the DA region. To maintain structural representations after they have been removed from the DA region, they must be established in LTM by more permanent associations—and indeed, more is needed than just associations, as I will explain next.

3.2. Structural Information in Long-Term Memory

So far, I have characterized LTM as an associative network of representations. But information in LTM must consist of more than just associations. LTM must be able to retain the specific relations between elements—for instance, who did what to whom when and where in an episode. One of the six requirements of a working memory system is that structural information can be learned over the long term. Associations are not enough to accomplish this. For instance, the fact that the pastor calmed the businessman cannot be represented by associating "pastor," "calm," and "businessman," because that

assembly could not be distinguished from one that represents ''The business-man calmed the pastor.'' The representation must specify the kind of relation between the concepts to distinguish different structures involving the same concepts. Thus, LTM must represent structures.

Does this imply that representations in LTM are structural representa-tions? Not necessarily—indeed, structural representations are an inconve-nient way of representing structures in LTM. A structural representation means that content components (i.e., objects, events) are linked to their contexts or roles by learned associations. For instance, ''pastor'' could be associated to the *agent* role in a proposition template, ''businessman'' to the *patient* role, and ''calming'' to the *action* role, to represent the fact that the pastor calmed the businessman. It is tempting to think of long-term learning in this way, because then all we needed is a mechanism that copies the bindings in the DA region into corresponding associations in LTM.

This scheme, however, would soon run into trouble when the system learns many more facts and events concerning the pastor and the businessman. Other events to be remembered could, for instance, involve the pastor in the patient role and the businessman in the agent role (e.g., the businessman bribing the pastor). Across all facts and events in memory, each object or concept would be associated to different roles, and each role to numerous objects and concepts, thus creating massive interference. To recover the fact that the pastor calmed the businessman (in one particular time and place), the system needs a mechanism to tell that the association of ''businessman'' with *patient* belongs together with the association of ''pastor'' to *agent* (as well as the association of information about time and place to the roles of *time* and *place*). In other words, there must be a mechanism to associate pair–wise content-role associations with each other. Because associations are not themselves representations, it is not obvious how they can be associated together. Therefore, long-term learning of structural information cannot simply consist of translating the bindings in WM into corresponding associations one-to-one.

3.2.1. Chunking Structural Representations

One solution to the problem of acquiring structural representations in LTM is to form unitized representations of structures, which I call *chunks* (Halford et al., 1998; Miller, 1956). A chunk is a representational unit in which other units and their relations are packed so that they are not individually accessi-ble—unless the chunk is unpacked again. Remembering structures as chunks requires that, for every new fact or episode, a new unitized repre-sentation must be created. Upon encountering for the first time an episode in which the pastor calms the businessman, the system must create a new chunk representing that proposition. The proposition chunk would be associated with the representations of the three concepts involved (i.e., pastor, calming, and businessman). In addition, it needs to be associated with representations that code the conjunction of each concept with its role

in the proposition (i.e., the conjunction of ''pastor'' with *agent*, of ''businessman'' with *patient*, and of ''calming'' with *action*). An illustration of how a relational representation in the DA region of WM is encoded as a chunk in LTM is given in Figure 5A (a similar structure of memory is proposed by Hummel & Holyoak, 2003).

There are several ways in which learning by chunking could be implemented in a neural network. With localist representations (i.e., each content component is represented by a single unit in the network), forming a new chunk means recruiting a new, as yet uncommitted unit to become the new chunk representation (Grossberg, 1987). With distributed representations (i.e., each content component is represented by a pattern of activation across many units), the new chunk could be represented by a new pattern created from its component patterns, for instance, by circular convolution (Jones & Mewhort, 2007; Plate, 2003).

3.2.2. Retrieval of Structural Information from LTM

Figure 5B illustrates how structural information from LTM can be retrieved back into the DA region of WM. The current contents of the DA region serve as retrieval cues for chunks in LTM. Using as cues only the content elements currently bound to some position or role in the DA region, however, would lead to confusion. For instance, when the DA region currently holds a mental model of a businessman calming a priest, the three content elements (priest, calming, and businessman) would cue a chunk representing the proposition ''the priest calmed the businessman'' and a chunk representing ''the businessman calmed the priest'' with equal strength. To focus retrieval better on those memories in LTM that match the current contents of thought with regard to their relations, the relations between concepts and roles in the DA region must act as additional cues. For this, the temporary bindings between content and role representations must first be translated into representations that code the conjunction of each concept and its current role. This is the same process as occurs during encoding of a structure as a chunk in LTM. One prediction from this assumption is that, when LTM is cued by a relational cue (e.g., a sentence or a pair of words), retrieval of memories matching the individual elements should be faster than retrieval of memories matching the relation between them, because the latter requires *ad hoc* chunking of the relational information in the cue as an intermediate step. This prediction is confirmed by several studies (e.g., Gronlund & Ratcliff, 1989; Gronlund, Edwards, & Ohrt, 1997; McElree, Dolan, & Jacoby, 1999).

To summarize, I assume that one important difference between the DA region and LTM is that in the DA region, structural information is represented explicitly, that is, by structural representations, whereas in LTM it is packed into chunks. Thus, although LTM contains information much richer than mere associations between ideas, its structure is purely associative—it is a network of associated chunks. Transferring structural
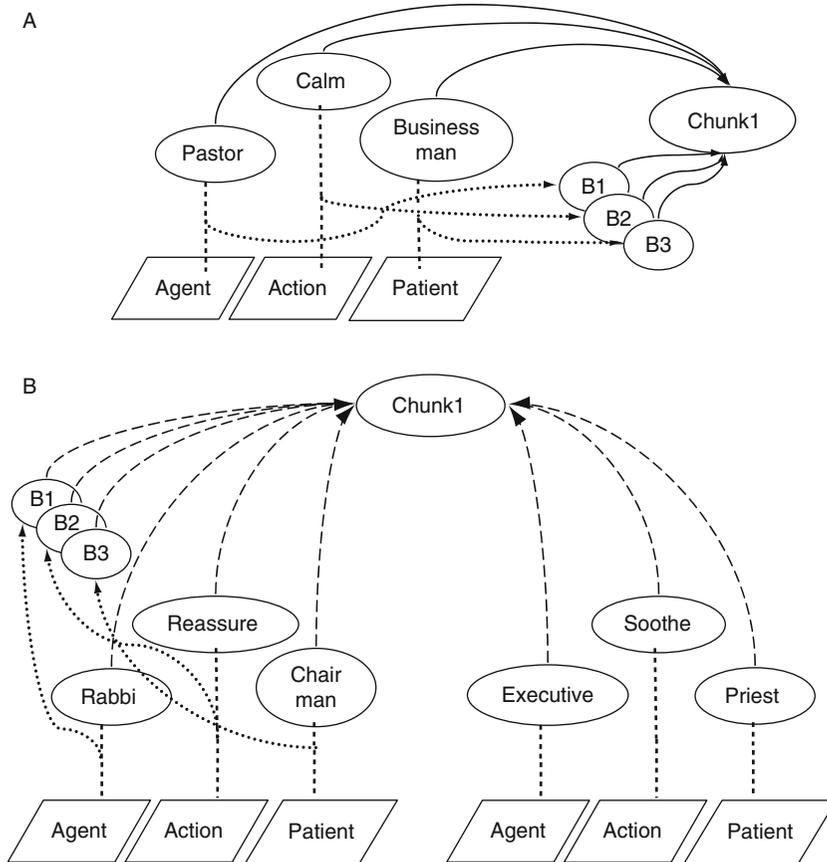
Figure 5 (A) The proposition ''The pastor calmed the businessman'' is represented in the direct-access region by binding the three concepts (pastor, calm, businessman) to their roles in a proposition template (agent, action, patient). Bindings are represented by thick dotted lines. To encode this proposition into LTM, a chunk is formed that includes information about the concepts, as well as about the bindings of each concept to its role (it is arbitrarily called *chunk1*). The bindings must first be translated into representations (B1 to B3). This conversion is represented by the thin-dotted lines. All representations packed into the chunk (pastor, calm, businessman, B1, B2, B3) are associated to that chunk; associations are shown as continuous lines. (B) When a new proposition is represented in the direct-access region, it cues related chunks in LTM for retrieval. A good analogy such as ''The rabbi reassured the chairman'' (illustrated on the left) cues *chunk1* with both its concepts and its bindings (again mediated through explicit representations of these bindings, B1 to B3). A bad analogy such as ''The executive soothed the priest'' (on the right) cues *chunk1* only by its concepts.

information from the DA region into LTM means packing the structure into a chunk. Retrieving of structural information from LTM means cueing a chunk with the current contents of the DA region, and their bindings, and

selecting the most strongly cued chunk to be unpacked into the DA region. Unpacking the chunk means to recreate the original structural representation in the DA region. This theory implies that relational information can be maintained in LTM, but it can be accessed for processing only after being projected back into the DA region. The DA region serves as a projection screen of our structural knowledge, thus enabling us to think about structures, that is, extract information about individual relations in a structure and manipulating individual components and relations in that structure.

3.2.3. Evidence for Chunking and Unpacking

Evidence for the mechanisms of encoding structural information into LTM by chunking, and retrieving it by cueing and unpacking, comes from three sources: analogical retrieval, the Hebb effect in list learning, and retrieval times for lists; I will discuss them in turn.

In analogical reasoning, people represent one structure in WM and try to retrieve some knowledge from LTM that matches the structure of the current representation in WM. For instance, when trying to figure out the physics of magma flowing down a volcano, the physics of water flowing down the drain could be a helpful analogue—although its elements differ, its causal structure is similar to the magma problem. Early research on the retrieval of analogues suggested that people are poor at finding structural matches in LTM. Instead, what mostly comes to mind seem to be memories that bear some surface similarity to the problem at hand, that is, similarity between the elements involved regardless of their structure. For instance, when reading a story, people are often reminded of other stories involving the same elements (e.g., the same animals) but rarely of stories that have an analogous causal structure but no shared elements (Gentner, Ratterman, & Forbus, 1993). Later research, however, led to a revision of this picture, showing that retrieval is sensitive to structural similarity between the cues and the to-be-retrieved memory. For instance, Wharton et al. (1994) found that, when given the sentence, ''The pastor calmed the businessman,'' people are more often reminded of a previously read sentence, ''The rabbi reassured the chairman'' than another previously read sentence, ''The executive soothed the priest.'' The two sentences in LTM have about equal similarity of their elements to the given sentence, but they differ in their structural similarity to the given sentence. Other research shows that analogues can be retrieved on the basis of structural similarity alone, without any similarity between the elements (Blanchette & Dunbar, 2000; Catrambone, 2002). Thus, there is evidence that retrieval from LTM is sensitive to the structure of the cue. This finding supports the assumption that not only the content elements currently held in the DA region, but also their bindings to their context, contribute to the cueing of chunks in LTM.

Evidence for chunking also comes from the literature on memory for lists. Lists can be thought of as structures in which a set of content elements

(i.e., list items) is linked to an ordered set of positions. One possible way of encoding lists into LTM is to build long-term associations between each item and its positions. When learning many different lists created by randomly reordering the same items—as in a typical experiment with many trials of immediate serial recall of digit lists—every item will eventually be associated with about equal strength to each position, making LTM virtually useless. However, if some items are placed in certain positions more often than in others, cumulative learning of these item-position associations would help remembering new lists in which the items are placed in these positions (Burgess & Hitch, 1999). This could be an explanation for the Hebb effect (Hebb, 1961): When across a series of trials of immediate serial recall the same list is repeated frequently (e.g., every third trial), performance on the repeated list selectively improves over trials. The explanation of the Hebb effect in terms of pair-wise long-term associations between items and positions, however, has received a serious blow by the finding that there is no cumulative learning when only every second item of a list is repeated in a constant position (Cumming, Page, & Norris, 2003). This and other findings (Hitch, Fastame, & Flude, 2005) have led Burgess and Hitch (2006) to revise their earlier account of the Hebb effect in terms of item-position associations. The revised model assumes instead that each list is remembered in LTM in a unitized form (i.e., a chunk). Repeated lists benefit from long-term knowledge to the degree that their initial items serve as retrieval cues for an existing list chunk in LTM. Thus, data and theorizing on serial recall converge with my analysis above on the conclusion that the representation of structures in LTM requires chunking.

Further evidence for the idea that structures are represented in LTM as unified chunks comes from retrieval times for lists from LTM. Several studies have compared the time for accessing memory lists from WM with the time for accessing the same lists from LTM. Access from LTM takes longer, but the additional time, which is assumed to reflect the time for retrieving a list from LTM into WM, does not depend on the length of the list (Conway & Engle, 1994; Oberauer, 2005b; Wickens, Moody, & Dow, 1981). This finding suggests that a list is retrieved as a single chunk, unpacking of which regenerates the list as a structure in the DA region. The unpacking process itself seems to proceed in parallel because otherwise, longer lists would take more time to unpack.[5]

---

[5] There are intriguing hints in the literature pointing to the existence of a simpler, purely associative learning mechanism alongside the mechanism for learning structures by chunking that I postulated above (Botvinick & Bylsma, 2005; Majerus, van der Linden, Mulder, Meulemans, & Peters, 2004). The associative learning system would gradually strengthen the association of all representations that cooccur in the DA region, thus picking up the probabilistic structure of the environment (insofar as it is reflected in the representation in the DA region). Exploring the role of this second learning system for the interaction of WM and LTM goes beyond the scope of this chapter.

3.3. Control of Encoding into and Retrieval from Long-Term Memory

Retrieval from LTM means that an LTM representation is projected into the DA region. Conversely, long-term learning means that the contents of the DA region are added to the representations in LTM. The WM system must have a mechanism to control the transfer between these two components. Research on incidental memory (Hebb, 1961; Hyde & Jenkins, 1969) and implicit memory (Graf & Schacter, 1985; Roediger, 1990) shows that every content of thought leaves a trace in LTM by default. Thus, translating the contents of the DA region into a corresponding trace in LTM seems to be the default mode of interaction between these two components.

Conversely, there is much evidence showing that long-term memories influence performance in immediate memory tasks. For instance, the Hebb effect mentioned earlier (Hebb, 1961) implies that memory traces of a repeated list support immediate recall of that list on further trials. Yet, it cannot be the case that any matching LTM content is automatically retrieved into the DA region, because if it was, retrieved information from activated LTM would flood the DA region. As a consequence, every new thought would immediately experience massive interference from many related old thoughts that it reminds us of. In fact, proactive interference is minimal in short-term memory or WM tasks (Davelaar, Goshen-Gottstein, Ashkenazi, Haarmann, & Usher, 2005; Halford, Maybery, & Bain, 1988; Oberauer & Vockenberg, 2009; Wickens, Born, & Allen, 1963). The WM system must have a mechanism that controls retrieval from LTM into the DA region to avoid proactive interference from LTM.

One solution, of course, is to shield the DA region against LTM by default, and open the gate for retrieval only on demand (i.e., triggered by a goal to retrieve long-term knowledge). This solution, however, would most of the time cut off ongoing cognitive processes from virtually all potentially relevant background knowledge. As a result, our mind would operate like a conventional computer that operates on a small set of representations in its current "work space," and is equipped with an entirely passive knowledge base, the contents of which must be retrieved by active search. Such a system encounters what has become known as the "frame problem" (Pylyshyn, 1987), that is, the problem of how to find, within reasonable time, the relevant information for the problem at hand in a vast data base.

Therefore, I think a more promising solution is to keep the gate between the DA region and LTM open by default, so that activated contents of LTM always have a chance to be retrieved and thereby to influence ongoing processing. Continuous interaction between the contents of the DA region and LTM does not, of course, in itself solve the frame problem; rather, it is

an architectural feature that arguably provides a better basis for a solution (cf. Shanahan & Baars, 2005): The contents of the DA region continually activate representations in LTM that match them or are otherwise associated with them, and activated representations in LTM continuously compete for being retrieved into the DA region. Thus, in most cases, the associative network of LTM takes part in the information processing loop to some degree, as I will elaborate in Section 4.

This assumption means that retrieval from LTM is "automatic" in the sense that it can occur without being intended. It is not clear yet whether retrieval is also "automatic" in the sense of not competing with other processes. Tasks testing retrieval from LTM have been found to compete with other concurrent cognitive operations in dual-task studies (Carrier & Pashler, 1995; Craik, Govoni, Naveh-Benjamin, & Anderson, 1996). In these studies, however, the retrieval task always involved an overt response reporting the retrieved information (e.g., saying aloud the word associated to a given cue, or making an old–new decision by pressing one of two keys in a recognition test). Therefore, it is difficult to tell from the existing evidence whether the locus of dual-task interference was the process of retrieving information from LTM into (declarative) WM, or the selection of a response to report that information. Because I assume that the response focus can hold only one response representation at a time, I would predict competition between overt retrieval and other concurrent tasks at response selection. This assumption could also explain why in dual-task studies of retrieval it is the secondary task that suffers, whereas retrieval is protected (Naveh-Benjamin, Craik, Perretta, & Toney, 2000). This would be expected if retrieval itself is not impaired by the secondary task at all, only its report disrupts or postpones the concurrent secondary task.[6]

Returning to the problem of controlling retrieval from LTM, a flexible solution is to put a threshold on the level of activation any LTM representation must attain to be retrieved into the DA region. That threshold could be adapted to the present processing goal—by lowering the threshold, the system could call for relevant knowledge in LTM, and by raising the threshold, it could ask not to be disturbed by intrusions from LTM. Adjusting this threshold to task demands must be regarded as one of the key executive functions regulating WM processes. In Section 4, I will discuss its role in striking the balance between an analytic and an associative mode of information processing.

---

[6] Carrier and Pashler (1995), using a variant of the PRP paradigm, argued that their findings cannot be explained by interference merely at the response selection stage because they found an additive effect of retrieval difficulty with their manipulation of stimulus–response asynchrony. I do not find this argument convincing, because it depends on the assumption that retrieval and response selection are distinct processing stages, and the difficulty of retrieval has no effect on the duration of response selection. I find it more plausible that their manipulation of retrieval difficulty affected the strength of the retrieved information, which translates into the rate at which evidence accumulates during response selection (Ratcliff, 1978).

4. Analytic and Associative Information Processing

The architecture of WM as outlined so far affords the implementation of two different modes of information processing. The central components of declarative WM—that is, the region of direct access and the focus of attention—together with the central components of procedural WM—that is, the bridge and the response focus—constitute what I call an *analytic processing subsystem*. It is complemented by an *associative subsystem* that consists of the activated part of LTM, including both declarative and procedural representations, which are not as strictly separated in (activated) LTM as they are in the central components of WM.

The distinction between analytic and associative kinds of processes has been noticed and discussed by several authors in the psychology of reasoning (Evans & Over, 1996; Sloman, 1996; Stanovich & West, 2000). Many of these theorists promote the distinction from one between processes to one between systems. The two systems go by various names, but the most common ones today are *System 1* and *System 2*. *System 1* is characterized as associative, intuitive, heuristic, fast, and automatic. *System 2* is thought of as analytic, deliberative, effortful, slow, controlled, and dependent on limited resources of WM (see Evans, 2008, for a review). I prefer to speak of analytic and associative processing modes, because I see them as the end points of a continuum, and they arise not from separate self-contained systems but from interacting sub-systems of WM.

The difference between the two processes or systems has been described by lists of contrasting features that are only loosely conceptually related. Here I make a more principled proposal based on the theory of WM outlined above. The analytic processing mode is characterized by three key features that distinguish it from the associative mode. First, it relies primarily—in the extreme case, exclusively—on declarative representations in the DA region, and procedural representations in the bridge, excluding influences from activated LTM. Second, it operates on representations with low dimensionality. Third, it enables independent selection and manipulation of declarative and procedural representations. I will elaborate on these three features below.

4.1. Shielding Central Working Memory Against Long-Term Memory

In the analytic mode, the central components of WM are shielded against the more peripheral component, that is, the activated part of LTM. This can be achieved by raising the threshold for retrieval from LTM, and by establishing strong bindings in the DA region and the bridge that cannot easily be replaced by activated LTM representations. On the declarative

side, these control measures ensure that processing draws only on a confined set of information held in the DA region. For instance, for a deductive inference, the reasoner should draw only on the information given in the premises, and refrain from taking his or her beliefs about the subject matter of the argument into account. People have difficulties narrowing their reasoning processes in this way—when asked whether a conclusion follows with logical validity from given premises, their judgments are strongly influenced by whether or not the conclusion matches their beliefs (for review see Klauer, Musch, & Naumer, 2000), and whether or not a counterexample to the conclusion is available in their knowledge base (De Neys, Schaeken, & d'Ydewalle, 2003; Markovits & Quinn, 2002). The degree to which deductive reasoning is shielded from knowledge in LTM is a function of intention and ability: The influence of knowledge can be reduced, but not completely eliminated, by strengthening the instruction to reason deductively (Vadeboncoeur & Markovits, 1999). Moreover, people with higher WM capacity are more successful in suppressing the influence of knowledge in cases where it conflicts with the logically correct answer (De Neys, 2006). These findings confirm the assumption that retrieval of knowledge associated with the current contents of the DA region is the default mode of processing, which needs to be actively prevented when the task demands exclusive reliance on a confined set of information represented in the DA region. Nevertheless, people can to some degree decouple their thinking from their factual knowledge to engage in counterfactual or hypothetical reasoning. Setting a high threshold for retrieval from LTM enables us to use the DA region as a blackboard for building new structural representations without suffering debilitating interference from potentially conflicting long-term memories or knowledge.

Similarly, on the procedural side, the analytic processing mode means that the selection of actions in response to perceived stimuli or to WM contents is determined by the task set in the bridge alone, which implements potentially arbitrary stimulus–response mappings, and suffers little or no bias from stimulus–response associations in LTM. Moreover, strong competing procedural representations in LTM cannot easily take over the bridge and wrest control from the currently implemented task set. Again, the degree to which processing is analytic can be controlled in response to task demands and characteristics of the environments. For instance, in the Stroop task, people can adjust the relative influence of two competing procedures, one supported by a strong LTM association (word reading) and one implementing the instructed task (color naming) in the bridge. When the proportion of incongruent trials is low, people allow a stronger contribution from word reading, compared to when the proportion of incongruent trials is high (Kane & Engle, 2003). As in the case of deductive reasoning, the degree to which processing is shielded from LTM influences depends not only on the executive processes that adjust the retrieval threshold according to

intentions and task demands, but also on the ability of the person to implement strong stimulus–response bindings in the bridge. For instance, in conditions with a low proportion of incongruent trials, people with low measured WM capacity suffered larger Stroop interference effects than those with high WM capacity (Kane & Engle, 2003).

4.2. Dimensionality of Representations

The assumption that analytic processing relies on representations with low dimensionality follows from my characterization of the DA region as providing a mental space in which representations of objects and events can be bound to positions. The dimensions of this mental space can be used to represent relations on any dimension, including physical space, time, and a host of feature dimensions, but the mental space is likely to have a limited number of dimensions—probably not more than two or three. Thus, the relations between elements in the DA region will be limited to a small number of dimensions. For instance, when the relative population sizes of China, Russia, and India are represented by placing the three countries on different points along a dimension of mental space, this representation selectively highlights one feature dimension, population size, while occluding all other features that we know these countries to have.

This contrasts with the representations of concepts, objects, and events in LTM, which can be characterized as points in a high-dimensional space (Jones & Mewhort, 2007; Landauer & Dumais, 1997), reflecting a wealth of knowledge by which we can distinguish them from each other. The dimensions of this "hyperspace" typically do not correspond to nameable feature dimensions. Projecting concepts, objects, or events into the region of direct access creates a low-dimensional representation in which a small subset of dimensions becomes explicitly represented by the dimensions of the mental coordinate system. Operating on representations in the DA region therefore implies focusing on a small number of feature dimensions, ignoring all other features of the represented entities. In contrast, processes that involve LTM representations directly draw on all feature knowledge linked to those representations in an unselective way.

The difference in dimensionality of representations entering analytic and associative processes is most evident in comparisons. Perceived stimuli, as well as representations held in the focus of attention, can be compared with representations in LTM directly, without having to retrieve the latter into the more central components of WM. Such a direct comparison or matching process is necessary for an efficient categorical identification of perceived objects and events: Everything we perceive is compared automatically to all representations in LTM, thus activating matching concepts as well as matching episodic representations. The same automatic comparison with LTM representations is also responsible for the generation

of the familiarity signal discussed above in the context of recognition. This comparison returns a global similarity or match value that sums across all feature dimensions for which values are available for both the perceived stimulus and the LTM representation it is compared to. The relative weight of each dimension can be modulated by cumulative learning but not by *ad hoc* task demands or intentions. Thus, the outcome of a comparison in purely associative processing mode is a global similarity value based on a weighted average across all available feature dimensions.

In analytic processing mode, in contrast, comparisons use representations of the entities to be compared in the DA region. Entities bound to places in mental space are compared by assessing their distance separately for each dimension; how these distance values are used is determined by the current task set. Thus, in the analytic mode, entities can be compared with regard to specific feature dimensions, ignoring all other dimensions, and a small number of dimensional distances can be assessed separately and combined in any arbitrary way. For instance, we can compare Russia and India simultaneously on two dimensions, size of population and geographical size, and judge that India has a larger population but Russia has a larger land area. We can also integrate both sizes by a weighted-average rule to judge which country is "larger." Alternatively, we might compute a ratio of the two values to judge which country is more densely populated.

A special case of comparison is comparing relations. For instance, we can compare the relationship between mother and daughter to that between grandmother and mother, or compare the relationship between the planets and the sun to the relationship between the electrons and the nucleus in the atom. Such relational comparisons underlie the *relational mapping* step in analogical reasoning (Gentner, 1989). Relational comparisons require access to individual relations on separable dimensions, rather than a global judgment of similarity or associatedness. These relations must be made explicit by projecting them into the region of direct access—for instance, the relation between planets and sun can be represented as a mental model in space in which smaller objects revolve around a larger central object (Gentner & Stevens, 1983). Thus, whereas the first step in analogical reasoning, finding an analogue, involves retrieval from LTM, the second step, relational mapping, relies exclusively on the region of direct access and thus constitutes a prime example of analytical processing. Therefore, analogical-mapping tasks, in which participants fill in a missing piece in a given analogy (of the form: A::B=C::?), can be regarded as relatively pure measures of the capacity of the DA region. Tasks of this kind are found in many tests of fluid intelligence. The Raven matrices, one of the best single measures of fluid intelligence (Marshalek, Lohman, & Snow, 1983), can be understood as analogical-mapping tasks with three times three elements (A::B::C=D::E::F=G::H::?). Fluid intelligence, in turn, is closely related to measures of WM capacity (Kane, Hambrick, & Conway, 2005; Oberauer,

Schulze, Wilhelm, & Süß, 2005). Together these facts support my assumption that measures of WM capacity essentially reflect the capacity of the DA region, which puts a limit on the complexity of structural representations that can be used in analytical reasoning.

The two modes of comparison—associative and analytical—map onto two corresponding modes of categorizing entities. Research on category learning has revealed two kinds of categorization tasks that are best solved by two distinct processing modes (Ashby & Maddox, 2005; Kloos & Sloutsky, 2008). Some categories are distinguished from each other by a single feature or a logical combination of very few features, and all other features of the entities to be categorized are irrelevant, and can vary widely within a category. Scientists and lawyers often strive to define concepts in this way—for instance, *iron* is defined by the number of protons and neutrons in its atoms, regardless of whether it appears in the shape of a nail or the hull of a ship. These categories are called *rule based* (Ashby & Maddox, 2005) because they can be described by relatively simple rules. Other categories are distinguished from each other on a large number of feature dimensions that combine in potentially complex, nonlinear ways (although linear combinations are easiest to learn, see Ashby & Maddox, 2005). The distinguishing features are often highly redundant, so that a random subset of them is sufficient for making highly accurate categorization decisions. These categories are called *information–integration* categories because categorization judgments are based on integrating feature information across many dimensions. Explicit categorization rules, if they can be formulated at all, are very complex because they must mention many features.

Kloos and Sloutsky (2008) capture essentially the same difference by their concept of *category density*: With dense categories, a large proportion of the features of exemplars are relevant for distinguishing members from nonmembers of the category, and the feature information is highly redundant, whereas with sparse categories, only one or a small subset of features is relevant, and all other features are uncorrelated with category membership. Therefore, learning of sparse categories is more dependent on selective attention to the relevant features. Kloos and Sloutsky show that sparse categories are best learned when the categorization rule is explicitly given, whereas dense categories are best learned when exemplars are presented. When given both kinds of information, adults learned rule-based representations of sparse categories, but similarity-based representations of dense categories. Children learned similarity-based representations for both kinds of categories, suggesting that forming similarity-based representations is the cognitively less demanding default approach to category learning. DeCaro, Thomas, and Beilock (2008) found that people with high WM capacity outperformed those with low capacity on learning low-dimensional rule-based categories, but performed worse than the low-capacity participants on learning higher dimensional information–integration categories.

Love (2002) provided evidence that higher dimensional categories, which are more difficult to learn than low-dimensional rule-based categories in a supervised learning regime, can be easier to learn in an unsupervised incidental learning regime.

Together these results support the distinction between two ways of learning categories, and two corresponding kinds of category representations. Dense categories can be represented as clusters in a high-dimensional feature space as provided by LTM. Their learning requires no selective attention to specific features, and new exemplars can be categorized on the basis of overall similarity (i.e., proximity to a cluster in the high-dimensional space). Sparse, rule-based categories can be represented as partitions in a low-dimensional mental space that includes only the relevant feature dimensions; the DA region provides a blackboard for defining that mental space. Exemplars are categorized by projecting their representations into that space, thus focusing exclusively on the relevant features. Learning of dense categories can occur incidentally, because all that is needed is the accumulation of exemplar memories in the high-dimensional feature space. Learning of sparse categories must rely on hypothesized (or given) rules that specify the relevant feature dimensions and the classification rule, and it requires engagement of the DA region of WM to represent the classification rule as a boundary in the space that is defined by the relevant feature dimensions. Therefore, learning of sparse, rule-based categories, but not of dense information-integration categories, depends on WM capacity.

4.3. Independent Selection of Declarative and Procedural Representations

I have argued above for separate sub-systems of declarative and procedural WM because I believe that this separation is an important feature for the functionality of the system. The separation preserves a high degree of independence between representations of objects and events on the one hand, and (cognitive) actions on the other hand.

Here I differentiate this claim, arguing that independence of declarative and procedural knowledge is a characteristic of the analytic subsystem. The DA region and the bridge are separate components, the contents of which are selected independently, and so are the focus of attention for declarative representations and the response focus for procedural representations. In activated LTM, however, I assume that declarative and procedural representations are more closely interwoven. Much evidence has been gathered showing that whenever a response is given to a stimulus, an association is built between the stimulus, the response, and the outcome, such that on the next occurrence of the same stimulus, the associated response and its outcome are primed (e.g., Hommel, 1998a; Rothermund, Wentura, & De Houwer, 2005). Learning such associations

is functional because it enables the cognitive system to acquire the probabilistic contingencies between stimuli (or, more generally, situations), responses, and outcomes. A large part of our everyday actions are routines executed in common environments—we encounter the same or very similar situations again and again, and usually respond to them in the same way as many times before, expecting and typically obtaining the same outcomes. Being able to rely on learned associations between situations, actions, and outcomes in such cases bypasses the need to make new decisions at every juncture. At the same time, the system must avoid being enslaved by learned routines. Even after we have driven the same way from home to work thousands of times, we want to be able to deviate from the common route toward a new destination. Thus, the system must have the choice between following the learned routine and responding to a situation in a new way. This is essentially the choice between the associative and the analytic processing mode.

In the associative mode, the perceived stimulus or situation activates a matching representation in declarative LTM. This declarative representation is associated to one or several responses to various degrees, based on previous experience, and these responses are therefore activated accordingly. In associative mode, the threshold for retrieval from LTM into the bridge is lowered, and there is no strong competing stimulus–response binding in the bridge. Therefore, the response most highly activated through LTM associations is retrieved and selected by the response focus. In analytic mode, in contrast, the retrieval threshold is set high, shielding the bridge from the habits and routines in procedural LTM, and the desired task set is implemented by strong stimulus–response bindings in the bridge. The perceived stimulus or situation still activates a matching representation in declarative LTM, which activates responses that are associated to that representation, but the activation of responses in LTM has only a relatively minor priming effect. Response selection is controlled by the task set in the bridge, and this task set is selected on the basis of the currently dominant goal, independent of the stimulus.

To summarize, the present theory takes on many ideas from dual-process or dual-system theories of reasoning (e.g., Evans, 2003; Sloman, 1996) and of action selection (e.g., Kornblum, Hasbroucq, & Osman, 1990) and integrates them in a common framework. I think of the two modes of processing not as outcomes of two separate, independent systems of information processing, but as the endpoints of a continuum. This continuum reflects the relative weight of the analytic subsystem of WM (i.e., the region of direct access and the bridge, together with the focus of attention and the response focus) and the associative subsystem (i.e., activated LTM representations and their associations). The place on this continuum is regulated by two parameters, the threshold for retrieval from activated LTM, and the strength of bindings in the central components of WM. The retrieval

threshold is set on a moment-to-moment basis by executive processes, and research on cognitive control is beginning to identify a number of variables that influence these executive processes, among them the amount of conflict between responses selected by the implemented task set and responses primed by associations (Botvinick, Braver, Barch, Carter, & Cohen, 2001), and the amount of performance pressure (Beilock & DeCaro, 2007). The second parameter, strength of bindings, is not under the control of executive processes but rather reflects the ability of the system to establish and maintain bindings and to minimize interference between them. This parameter therefore reflects a relatively stable characteristic of an individual's WM system, and it is the main source of individual differences in WM capacity and performance on tasks requiring analytical processing.

5. Concluding Remarks

Readers will have noticed that the "working memory" sketched here is not genuinely a memory. Rather, it is an attentional system that interacts equally with perception and with (long-term) memory. This can be illustrated with an example investigated by Halford and his colleagues, the comprehension of statistical interactions (Halford, Baker, McCredden, & Bain, 2004). The necessary information can be fully available for perception (e.g., in the form of a table or a bar chart) and yet people have severe difficulties grasping a three-way interaction and generally fail with a four-way interaction. A similar argument can be made with respect to many intelligence test tasks used to measure reasoning ability (Oberauer, Süß, Wilhelm, & Wittmann, 2003). The limiting factor for complex reasoning is not our ability to remember all the relevant pieces of information but to put them together by binding them into a common schema or a common cognitive coordinate system. This capacity for *relational integration* (Robin & Holyoak, 1995) applies equally to representations from memory and to representation of the perceived environment. My colleagues and I have developed tasks that measure WM capacity by testing people's ability to integrate given information into structures. The given information was either continually visible or had to be maintained in memory. Both task versions were highly correlated and were equally good predictors of reasoning ability (Oberauer, Süß, Wilhelm, & Wittmann, 2008).

WM has not evolved to serve us as a short-term store. The fact that it can also be used as such is better understood as a by-product of its ability to uphold a structural representation and shield it to some degree from unwanted, potentially interfering input from both perception and LTM. The main function of WM is to serve as a blackboard for information processing on which we can construct new representations with little

interference from old memories, knowledge, and perceptual input, thus enabling us to investigate a hypothetical alternative state of reality, a future state—or some aspect of the past.

ACKNOWLEDGMENTS

REFERENCES

Allport, A. (1987). Selection for action: Some behavioral and neurophysiological considerations of attention and action. In H. Heuer & A. F. Sanders (Eds.), *Perspectives on perception and action* (pp. 395–419). Hillsdale: Erlbaum.

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge: Harvard University Press.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–178.

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

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *Recent advances in learning and motivation Vol. VIII*, (pp. 47–90). New York: Academic Press.

Bar-Anan, Y., Liberman, N., Trope, Y., & Algom, D. (2007). Automatic processing of psychological distance: Evidence from a Stroop task. *Journal of Experimental Psychology: General*, *136*, 610–622.

Beilock, S. L., & DeCaro, M. S. (2007). From poor performance to success under stress: Working memory, strategy selection, and mathematical problem solving under pressure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 983–998.

Bertelson, P. (1965). Serial choice reaction-time as a function of response versus signal-and-response repetition. *Nature*, *2006*, 217–218.

Blanchette, I., & Dunbar, K. (2000). How analogies are generated: The roles of structural and superficial similarity. *Memory & Cognition*, *28*, 108–124.

Blankenship, A. B. (1938). Memory span: A review of the literature. *Psychological Bulletin*, *35*, 1–25.

Bogacz, R., & Gurney, K. (2008). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation*, *19*, 442–477.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624–652.

Botvinick, M. M., & Bylsma, L. M. (2005). Regularization in short-term memory for serial order. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *31*, 351–358.

Brambosch, A. (2003). *Auf der Suche nach den Grenzen humaner kognitiver Simultanverarbeitung*. Unpublished Diploma thesis, Universität Potsdam, Potsdam.

Bryck, R. L., & Mayr, U. (2005). On the role of verbalization during task set selection: Switching or serial order control? *Memory & Cognition*, *33*, 611–623.

Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, *106*, 551–581.

Burgess, N., & Hitch, G. J. (2005). Computational models of working memory: Putting long-term memory into context. *Trends in Cognitive Sciences*, *9*, 535–541.

Burgess, N., & Hitch, G. J. (2006). A revised model of short-term memory and long-term learning of verbal sequences. *Journal of Memory and Language*, *55*, 627–652.

Byrne, M. D., & Anderson, J. R. (2001). Serial modules in parallel: The psychological refractory period and perfect time-sharing. *Psychological Review*, *108*, 847–869.

Carrier, L. M., & Pashler, H. (1995). Attentional limits in memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1339–1348.

Catrambone, R. (2002). The effects of surface and structural feature matches on the access of story analogs. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *28*, 318–334.

Cohen-Kdoshay, O., & Meiran, N. (2007). The representation of instructions in working memory leads to autonomous response activation: Evidence from the first trials in the flanker paradigm. *Quarterly Journal of Experimental Psychology*, *60*, 1140–1154.

Conway, A. R. A., & Engle, R. W. (1994). Working memory and retrieval: A resource-dependent inhibition model. *Journal of Experimental Psychology: General*, *123*, 354–373.

Conway, A. R. A., Kane, M. J., & Engle, R. W. (2003). Working memory capacity and its relation to general intelligence. *Trends in Cognitive Sciences*, *7*, 547–552.

Cowan, N. (1988). Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin*, *104*, 163–191.

Craik, F. I. M., Govoni, R., Naveh-Benjamin, M., & Anderson, N. D. (1996). The effects of divided attention on encoding and retrieval processes in human memory. *Journal of Experimental Psychology: General*, *125*, 159–180.

Cumming, N., Page, M., & Norris, D. (2003). Testing a positional model of the Hebb effect. *Memory*, *11*, 43–63.

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

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

Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J., & Usher, M. (2005). The demise of short-term memory revisited: Empirical and computational investigation of recency effects. *Psychological Review*, *112*, 3–42.

De Neys, W. (2006). Dual processing in reasoning. Two systems but one reasoner. *Psychological Science*, *17*, 428–433.

De Neys, W., Schaeken, W., & d'Ydewalle, G. (2003). Inference suppression and semantic memory retrieval: Every counterexample counts. *Memory & Cognition*, *31*, 581–595.

DeCaro, M. S., Thomas, R. D., & Beilock, S. L. (2008). Individual differences in category learning: Sometimes less working memory capacity is better than more. *Cognition*, *107*, 284–294.

DeSoto, C., London, M., & Handel, S. (1965). Social reasoning and spatial paralogic. *Journal of Personality and Social Psychology*, *2*, 513–521.

Dreisbach, G., & Goschke, T. (2004). How positive affect modulates cognitive control: Reduced perseveration at the cost of increased distractability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 343–353.

Duncan, J., Emslie, H., Williams, P., Johnson, R., & Freer, C. (1996). Intelligence and the frontal lobe: The organization of goal-directed behavior. *Cognitive Psychology*, *30*, 257–303.

Duncan, J., Parr, A., Woolgar, A., Thompson, R., Bright, P., Cox, S., et al. (2008). Goal neglect and Spearman's g: Competing parts of a complex task. *Journal of Experimental Psychology: General*, *137*, 131–148.

Durstewitz, D., & Seamans, J. K. (2002). The computational role of dopamine D1 receptors in working memory. *Neural Networks*, *15*, 561–572.

Ellenbogen, R., & Meiran, N. (2008). Working memory involvement in dual taks performance: Evidence from the backward compatibility effect. *Memory & Cognition, 36*, 968–978.

Emerson, M. J., & Miyake, A. (2003). The role of inner speech in task switching: A dual-task investigation. *Journal of Memory and Language, 48*, 148–168.

Evans, J. S. B. T. (2003). In two minds: Dual process accounts of reasoning. *Trends in Cognitive Sciences, 10*, 454–459.

Evans, J. S. B. T. (2008). Dual-process accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology, 59*, 6.1–6.24.

Evans, J. S. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Hove: Psychology Press.

Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin & Review, 9*, 59–79.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28*, 3–71.

Friedman, N. P., & Miyake, A. (2004). The relations among inhibition and interference control functions: A latent variable analysis. *Journal of Experimental Psychology: General, 133*, 101–131.

Friedman, N. P., Miyake, A., Corley, R. P., Young, S. E., DeFries, J. C., & Hewitt, J. K. (2006). Not all executive functions are related to intelligence. *Psychological Science, 17*, 172–179.

Friedman, N. P., Miyake, A., Young, S. E., DeFries, J. C., Corley, R. P., & Hewitt, J. K. (2008). Individual differences in executive functions are almost entirely genetic in origin. *Journal of Experimental Psychology: General, 137*, 201–225.

Garavan, H. (1998). Serial attention within working memory. *Memory & Cognition, 26*, 263–276.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). Cambridge: Cambridge University Press.

Gentner, D., Ratterman, M. J., & Forbus, K. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology, 25*, 524–575.

Gentner, D., & Stevens, A. L. (1983). *Mental Models*. Hillsdale: Erlbaum.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91*, 1–67.

Goldman-Rakic, P. (1997). Architecture of the prefrontal cortex and the central executive. *Annals of the New York Academy of Science, 769*, 71–83.

Goodwin, G. P., & Johnson-Laird, P. N. (2005). Reasoning about relations. *Psychological Review, 112*, 468–493.

Göthe, K., Oberauer, K., & Kliegl, R. (2007). Age differences in dual-task performance after practice. *Psychology and Aging, 22*, 596–606.

Graf, P., & Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, Memory & Cognition, 11*, 501–518.

Gronlund, S. D., Edwards, M. B., & Ohrt, D. D. (1997). Comparison of the retrieval of item versus spatial position information. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 1261–1274.

Gronlund, S. D., & Ratcliff, R. (1989). Time course of item and associative information: Implication for global memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 846–858.

Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science, 11*, 23–63.

Grossberg, S., & Stone, G. (1986). Neuronal dynamics of attention switching and temporal order information in short term memory. *Memory & Cognition, 14*, 451–468.

Halford, G. S., Baker, R., McCredden, J. E., & Bain, J. D. (2004). How many variables can humans process? *Psychological Science, 16*, 70–76.

Halford, G. S., Maybery, M. T., & Bain, J. D. (1988). Set-size effects in primary memory: An age-related capacity limitation? *Memory & Cognition*, *16*, 480–487.

Halford, G. S., Wilson, W. H., & Phillips, S. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences*, *21*, 803–864.

Hebb, D. O. (1961). Distinctive features of learning in the higher animal. In J. F. Delafresnaye (Ed.), *Brain mechanisms and learning* (pp. 37–46). Oxford: Blackwell.

Hitch, G. J., Fastame, M. C., & Flude, B. (2005). How is the serial order of a verbal sequence coded? Some comparisons between models. *Memory*, *13*, 247–258.

Hommel, B. (1998a). Event files: Evidence for automatic integration of stimulus–response episodes. *Visual Cognition*, *5*, 183–216.

Hommel, B. (1998b). The prepared reflex: Automaticity and control in stimulus–response translation. In S. Monsell & J. Driver (Eds.). *Attention and Performance XVIII: Control of Cognitive Processes*. Cambridge, MA: MIT Press.

Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, *24*, 849–937.

Hübner, M., Kluwe, R. H., Luna-Rodriguez, A., & Peters, A. (2004). Response selection difficulty and asymmetrical costs of switching between tasks and stimuli: No evidence for an exogenous component of task-set reconfiguration. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 1043–1063.

Hummel, J. E., & Holyoak, J. A. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220–264.

Hyde, T. S., & Jenkins, J. J. (1969). Differential effects of incidental tasks on the organization of recall of a list of highly associated words. *Journal of Experimental Psychology*, *82*, 472–481.

Jentzsch, I., & Dudschig, C. (2009). Why do we slow down after an error? Mechanisms underlying the effect of posterror slowing. *Quarterly Journal of Experimental Psychology*, *62*, 209–218.

Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. Hillsdale: Erlbaum.

Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, *114*, 1–38.

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

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, *24*, 175–219.

Kane, M. J., & Engle, R. W. (2003). Working-memory capacity and the control of attention: The contributions of goal neglect, response competition, and task set to Stroop interference. *Journal of Experimental Psychology: General*, *132,* (47–70).

Kane, M. J., Hambrick, D. Z., & Conway, A. R. A. (2005). Working memory capacity and fluid intelligence are strongly related constructs: Comment on Ackerman, Beier, and Boyle (2004). *Psychological Bulletin*, *131*, 66–71.

Kiesel, A., Wendt, M., & Peters, A. (2007). Task switching: On the origin of response congruency effects. *Psychological Research*, *71*, 117–125.

Klauer, K. C., Musch, J., & Naumer, B. (2000). On belief bias in syllogistic reasoning. *Psychological Review*, *107*, 852–884.

Klauer, K. J. (1996). Teaching inductive reasoning: Some theory and three experimental studies. *Learning and Instruction*, *6*, 37–57.

Kleinsorge, T. (1999). Response repetition benefits and costs. *Acta Psychologica*, *103*, 295–310.

Kloos, H., & Sloutsky, V. M. (2008). What's behind different kinds of kind: Effects of statistical density on learning and representation of categories. *Journal of Experimental Psychology: General*, *137*, 52–72.

Knauff, M., Mulack, T., Kassubek, J., Salih, H. R., & Greenlee, M. W. (2002). Spatial imagery in deductive reasoning: A functional MRI study. *Cognitive Brain Research*, *13*, 203–212.

Kornblum, S., Hasbroucq, T., & Osman, A. M. (1990). Dimensional overlap: Cognitive basis for stimulus–response-compatibility—A model and taxonomy. *Psychological Review*, *97*, 253–270.

Kray, J., Li, K. Z. H., & Lindenberger, U. (2002). Age differences in executive functioning: The search for sources of age-differential decline in task-switching components. *Brain & Cognition*, *49*, 363–381.

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

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.

Laming, D. (1968). *Information theory of choice reaction times*. London: Academic Press.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.

Lewandowsky, S. (1999). Redintegration and response suppression in serial recall: A dynamic network model. *International Journal of Psychology (Special Issue: Short-Term/Working Memory)*, *34*, 434–446.

Lien, M.-C., & Proctor, R. W. (2002). Stimulus-response compatibility and psychological refractory period effects: Implications for response selection. *Psychonomic Bulletin & Review*, *9*, 212–238.

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*, 295–327.

Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, *108*, 393–434.

Love, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin & Review*, *9*, 829–835.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, *109*, 163–203.

Majerus, S., van der Linden, M., Mulder, L., Meulemans, T., & Peters, F. (2004). Verbal short-term memory reflects the sublexical organization of the phonological language network: Evidence from an incidental phonotactic learning paradigm. *Journal of Memory and Language*, *51*, 297–306.

Markovits, H., & Quinn, S. (2002). Efficiency of retrieval correlates with "logical" reasoning from causal conditional premises. *Memory & Cognition*, *30*, 696–706.

Marr, D. (1982). *Vision*. San Francisco: Freeman.

Marshalek, B., Lohman, D. F., & Snow, R. E. (1983). The complexity continuum in the radex and hierarchical models of intelligence. *Intelligence*, 7, 107–127.

Mayr, U. (2001). Age differences in the selection of mental sets: The role of inhibition, stimulus ambiguity, and response-set overlap. *Psychology & Aging*, *16*, 96–109.

Mayr, U., & Kliegl, R. (2000). Task-set switching and long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 1124–1140.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.

McElree, B., Dolan, P. O., & Jacoby, L. L. (1999). Isolating the contributions of familiarity and source information to item recognition: A time course analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 563–582.

McKone, E. (1998). The decay of short-term implicit memory: Unpacking lag. *Memory & Cognition*, *26*, 1173–1186.

Meiran, N., Chorev, Z., & Sapir, A. (2000). Component processes in task-switching. *Cognitive Psychology*, *41*, 211–253.

Meiran, N., & Kessler, Y. (2008). The task-rule congruency effect in task switching reflects activated long-term memory. *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 137–157.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.

Mitchell, D. B. (2006). Nonconscious priming after 17 years. *Psychological Science*, *17*, 925–929.

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex ''frontal lobe'' tasks: A latent variable analysis. *Cognitive Psychology*, *41*, 49–100.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7, 134–140.

Nairne, J. S. (1990). A feature model of immediate memory. *Memory & Cognition*, *18*, 251–269.

Naveh-Benjamin, M., Craik, F. I. M., Perretta, J. G., & Toney, S. T. (2000). The effects of divided attention on encoding and retrieval processes: The resiliency of retrieval processes. *Quarterly Journal of Experimental Psychology*, *53A*, 609–625.

Nelson, J. K., Reuter-Lorenz, P., Sylvester, C.-Y. C., Jonides, J., & Smith, E. E. (2003). Dissociable neural mechanisms underlying response-based and familiarity-based conflict in working memory. *Proceedings of the National Academy of Sciences*, *100*, 11171–11175.

Newell, A. (1990). *Unified theories of cognition*. Cambridge: Harvard University Press.

Nipher, F. E. (1878). On the distribution of errors in numbers written from memory. *Transactions of the Academy of Science of St Louis*, *3*, CCX–CCXI.

Norman, D. A., & Shallice, T. (1980). *Attention to action. Willed and automatic control of behavior*. (CHIP Report No. 99) San Diego: University of California.

Oberauer, K. (2001). Removing irrelevant information from working memory. A cognitive aging study with the modified Sternberg task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 948–957.

Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 411–421.

Oberauer, K. (2003). Selective attention to elements in working memory. *Experimental Psychology*, *50*, 257–269.

Oberauer, K. (2005a). Binding and inhibition in working memory—Individual and age differences in short-term recognition. *Journal of Experimental Psychology: General*, *134*, 368–387.

Oberauer, K. (2005b). Control of the contents of working memory—A comparison of two paradigms and two age groups. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 714–728.

Oberauer, K. (2005c). The measurement of working memory capacity. In O. Wilhelm & R. W. Engle (Eds.), *Handbook of understanding and measuring intelligence* (pp. 393–408). Thousand Oaks: Sage.

Oberauer, K. (2008). How to say no: Single- and dual-process theories of short-term recognition tested on negative probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 439–459.

Oberauer, K., & Göthe, K. (2006). Dual-task effects in working memory: Interference between two processing tasks, between two memory demands, and between storage and processing. *European Journal of Cognitive Psychology*, *18*, 493–519.

Oberauer, K., & Kliegl, R. (2001). Beyond resources: Formal models of complexity effects and age differences in working memory. *European Journal of Cognitive Psychology*, *13*, 187–215.

Oberauer, K., & Kliegl, R. (2004). Simultaneous execution of two cognitive operations—Evidence from a continuous updating paradigm. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 689–707.

Oberauer, K., & Kliegl, R. (2006). A formal model of capacity limits in working memory. *Journal of Memory and Language*, *55*, 601–626.

Oberauer, K., & Lange, E. B. (2008). Interference in verbal working memory: Distinguishing similarity-based confusion, feature overwriting, and feature migration. *Journal of Memory and Language*, *58*, 730–745.

Oberauer, K., & Lange, E. B. (2009). Activation and binding in verbal working memory: A dual-process model for the recognition of nonwords. *Cognitive Psychology*, *58*, 102–136.

Oberauer, K., & Lewandowsky, S. (2008). Forgetting in immediate serial recall: Decay, temporal distinctiveness, or interference? *Psychological Review*, *115*, 544–576.

Oberauer, K., Schulze, R., Wilhelm, O., & Süß, H.-M. (2005). Working memory and intelligence—Their correlation and their relation: A comment on Ackerman, Beier, and Boyle (2005). *Psychological Bulletin*, *131*, 61–65.

Oberauer, K., Süß, H.-M., Wilhelm, O., & Sander, N. (2007). Individual differences in working memory capacity and reasoning ability. In A. R. A. Conway, C. Jarrold, M. J. Kane, A. Miyake, & J. N. Towse (Eds.), *Variation in working memory* (pp. 49–75). New York: Oxford University Press.

Oberauer, K., Süß, H.-M., Wilhelm, O., & Wittmann, W. W. (2003). The multiple faces of working memory—Storage, processing, supervision, and coordination. *Intelligence*, *31*, 167–193.

Oberauer, K., Süß, H.-M., Wilhelm, O., & Wittmann, W. W. (2008). Which working memory functions predict intelligence? *Intelligence*, *36*, 641–652.

Oberauer, K., & Vockenberg, K. (2009). Updating of working memory: Lingering bindings. *Quarterly Journal of Experimental Psychology*, *62*, 967–987.

Page, M. P. A., & Norris, D. (1998). The primacy model: A new model of immediate serial recall. *Psychological Review*, *105*, 761–781.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, *116*, 220–244.

Plate, T. A. (2003). Convolution-based memory models. In L. Nadel (Ed.), *Encyclopedia of cognitive science* (pp. 824–828). London: Nature Publishing Group.

Pylyshyn, Z. W. (1987). *The robot's dilemma. The frame problem in Artificial Intelligence*. Norwood: Ablex.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 50–109.

Robin, N., & Holyoak, K. J. (1995). Relational complexity and the functions of prefrontal cortex. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 987–997). Cambridge: MIT Press.

Roediger, H. L. (1990). Implicit memory: Retention without remembering. *American Psychologist*, *45*, 1043–1056.

Roelofs, A. (2001). Set size and repetition matter: Comment on Caramazza and Costa (2000). *Cognition*, *80*, 283–290.

Rothermund, K., Wentura, D., & De Houwer, J. (2005). Retrieval of incidental stimulus-response associations as a source of negative priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 482–495.

Rubin, O., & Meiran, N. (2005). On the origins of the task mixing cost in the cuing task-switching paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 1477–1491.

Schmiedek, F., Oberauer, K., Wilhelm, O., Süß, H. M., & Wittmann, W. W. (2007). Individual differences in components of reaction time distributions and their relations to working memory and intelligence. *Journal of Experimental Psychology: General*, *136*, 414–429.

Schneider, D. W., & Logan, G. D. (2007). Retrieving information from a hierarchical plan. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 1076–1091.

Schuch, S., & Koch, I. (2004). The costs of changing the representation of action: Response repetition and response-response compatibility in dual tasks. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 566–582.

Shanahan, M., & Baars, B. J. (2005). Applying global workspace theory to the frame problem. *Cognition*, *98*, 157–176.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3–22.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, *23*, 645–726.

Süß, H.-M., Oberauer, K., Wittmann, W. W., Wilhelm, O., & Schulze, R. (2002). Working memory capacity explains reasoning ability—And a little bit more. *Intelligence*, *30*, 261–288.

Unsworth, N., & Engle, R. E. (2007). The nature of individual differences in working memory capacity: Active maintenance in primary memory and controlled search from secondary memory. *Psychological Review*, *114*, 104–132.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.

Vadeboncoeur, I., & Markovits, H. (1999). The effect of instructions and information retrieval on accepting the premises in a conditional reasoning task. *Thinking & Reasoning*, *5*, 97–113.

Waszak, F., Hommel, B., & Allport, A. (2003). Task-switching and long-term priming: Role of episodic stimulus-task bindings in task-shift costs. *Cognitive Psychology*, *46*, 361–413.

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., Wickens, T. D., & Melz, E. R. (1994). Below the surface: Analogical similarity and retrieval competition in reminding. *Cognitive Psychology*, *26*, 64–101.

Whittlesea, B. W. A., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language*, *29*, 716–732.

Wickens, D. D., Born, D. G., & Allen, C. K. (1963). Proactive inhibition and item similarity in short-term memory. *Journal of Verbal Learning & Verbal Behavior*, *2*, 440–445.

Wickens, D. D., Moody, M. J., & Dow, R. (1981). The nature and timing of the retrieval process and of interference effects. *Journal of Experimental Psychology: General*, *110*, 1–20.

Wilhelm, O., & Oberauer, K. (2006). Why are reasoning ability and working memory capacity related to mental speed? An investigation of stimulus–response compatibility in choice–reaction–time tasks. *European Journal of Cognitive Psychology*, *18*, 18–50.

Woltz, D. J., & Was, C. A. (2007). Available but unattended conceptual information in working memory: Temporarily active semantic content or persistent memory for prior operations? *Journal of Experimental Psychology: Learning, Memory & Cognition*, *33*, 155–168.