2010

Spontaneous remembering is the norm: What integrative models tell us about human consciousness and memory

S. Franklin

B. Baars

Chapter 6

Title: Spontaneous remembering is the norm: What integrative models tell us about human consciousness and memory.

Author(s): Stan Franklin & Bernard J. Baars

Running head: Spontaneous remembering is the norm

To appear in J. H. Mace (Ed.): The act of remembering: Toward an understanding of how we recall the past. Wiley-Blackwell

Introduction

Ever since Hermann Ebbinghaus, the scientific study of memory has focused on deliberate memorizing and recall in laboratory experiments. However, “memorizing,” while experimentally convenient, is rather uncommon in everyday life (Rubin 2006). Based on some five decades of thought-monitoring studies, we know that most of our normal, spontaneous thoughts do not involve explicit recall of novel, deliberately memorized material. Rather, the spontaneous stream keeps coming back to our "current concerns," to answer the question, *"What do I do next, to reach my most important goals?"* The stream of thought reveals a wide range of spontaneous thoughts, perceptual experiences, unbidden memories, fantasies and feelings, reveries, emotionally toned fringe experiences, feelings of effort, familiarity and unfamiliarity, and self-evaluating thoughts (James 1890; Chafe 2000; Schooler, 2002; Epstein 2000). We can acquire memories simply by paying attention to novel and important events, with no explicit attempt to memorize them, and we constantly make use of our memory systems while acting, thinking, and coping with the challenges of everyday life. Real-life memory helps to solve real-life problems.

Here we explore some well-established features of the stream of consciousness: spontaneous remembering, including incidental learning and problem solving, expectation-driven recall, action control, and the “availability heuristic” --- the influence of conscious accessibility on judgments and decisions spontaneous recall and problem-solving (Tversky & Kahneman, 1973). These empirical phenomena emerge naturally from Global Workspace Theory as implemented in the LIDA model (see Baars, 1988, 1997, 2002; Franklin, 2001, Ramamurthy, D'Mello and Franklin, 2006), an integrative, evidence-based computational model of cognition. LIDA has a detailed role for both conscious and unconscious processes, based on the Global Workspace Theory (Franklin 2001; Ramamurthy, D'Mello, and Franklin 2006; Baars 1988, 2002). It reveals an adaptive role for the stream of consciousness, interacting with well-established memory systems and with the external world, making constant use of declarative, perceptual, transient episodic, and procedural memory types.

A neglected scientific question is, "How do the human memory systems enable spontaneous, life-relevant retrieval of everyday information?" We explore a set of spontaneous memory retrieval phenomena, with explanations based on the LIDA-GWT model of conscious and unconscious goal-directed cognition. Following are some features of natural thinking that require explanation. Our basic message is *that such spontaneous* memories and thoughts comprise the normal, everyday stream of thoughts (James 1890 Chapter 9; Chafe 2000; Epstein 2000; Smallwood et al, 2004). This view suggests that spontaneous thoughts are not just irrelevant "mind-wandering," as some researchers suggest, but rather a highly functional, implicitly purposeful, problem-solving stream. The function of spontaneous thoughts is to solve life problems, including implicit ones.

1. Spontaneous thinking involves implicit problem-solving.

Baars (1988, 1997) has made the case that the fundamental unit of spontaneous thought is a "C-U-C" triplet, consisting of a conscious, then an unconscious, and finally a conscious stage of thought. The basic notion is that the apparently accidental quality of the free flow of conscious thoughts results from many intertwining C-U-C problem-solving threads. Notice that these triplets are not typically *labeled* as "problem-solving". They are typically shaped and guided by unconscious contextual expectations and goals, and therefore they are not self-consciously labeled or identified at all in a metacognitive fashion. The trick is that apparently irrelevant thoughts, or "mind-wandering", may actually be highly purposeful, life-relevant problem-solving.

These three proposed stages of spontaneous problem-solving, which are well-known in numerous experimental problem-solving tasks, are:

(1) C1. *Conscious* problem identification;

Any conscious "prime" creates expectations. That may be true for unconscious events as well. In language perception, for example, a normal sentence can be stopped in the middle, and listeners will still be able to make highly reliable predictions about the next word, and certainly about the syntactic category of that word. (So-called Cloze sentences.) Much the same can be said about watching a football game or perceiving any other structured activity, like listening to a song. In general, conscious events evoke numerous expectations.

Notice that when we encounter a missing word like ____, we do not necessarily tell ourselves explicitly to bring it to mind. The answer simply tends to emerge in consciousness spontaneously, much like the classical Gestalt closure phenomena. Thus the brain spontaneously performs implicit problem-solving *as if* it were trying to identify and answer questions.

The same thing is likely to be true of endogenous conscious events, like spontaneous thoughts, memories, images and feelings. In general, conscious primes can be considered to *present a problem* for the nervous system to solve, involving a set of predictions about what is coming next.

The next stage in problem-solving is often (though not always) unconscious.

(2) U. *Unconscious* incubation;

Going back to Gestalt psychology a number of researchers have studied unconscious or implicit problem-solving. These stages of problem-solving are traditionally called "incubation". We see incubation with the tip of the tongue phenomenon, which has been described since William James as a silent (i.e., unconscious) anticipation of the word whose meaning we know, but whose phonological form we cannot bring to mind. Such unconscious states of expectation appear to be active, in the sense that the answer to the question – e.g., *"what is the name of the flying dinosaurs?"* – is being actively pursued. Thus unconscious incubation is not merely a passive waiting for the answer to appear, but seems to involve active problem-solving.

Finally, the third stage is conscious.

(3) C2. *Conscious* emergence of an answer.

In the case of the tip of the tongue phenomenon we expect the answer to emerge consciously. But conscious emergence of an answer applies to a dizzying variety of problem-solving tasks. Associative memory is certainly one of those, simply because an established association allows us to present associate *A* consciously, and then expect associate *B* to emerge in consciousness. Pattern recognition and action planning show the same regularity.

Very many psychological tasks reveal a Conscious-Unconscious-Conscious format, ranging from ambiguity resolution in language and perception, to visual target search, question-answering, free association, decision-making, action control, path navigation and word retrieval. Each of those well-studied tasks begin with a conscious phase, a limited period of forgetting or distraction, and the spontaneous appearance of a conscious answer. In Global Workspace Theory (GWT) the conscious moments allow collaborative interactions to occur among multiple unconscious knowledge sources, using a shared momentary memory domain called a global workspace (Baars, 1988, 2002). In the brain, unconscious knowledge sources may involve declarative or procedural memory, primed neuronal networks, or all those systems working collaboratively.

When answers are found by the automatic routines, those answers may become conscious again (C2), in a process that was traditionally called “insight.” However, no profound or surprising insights are required to observe these phenomena. C1-U-C2 patterns seems to happen every minute of the waking day.

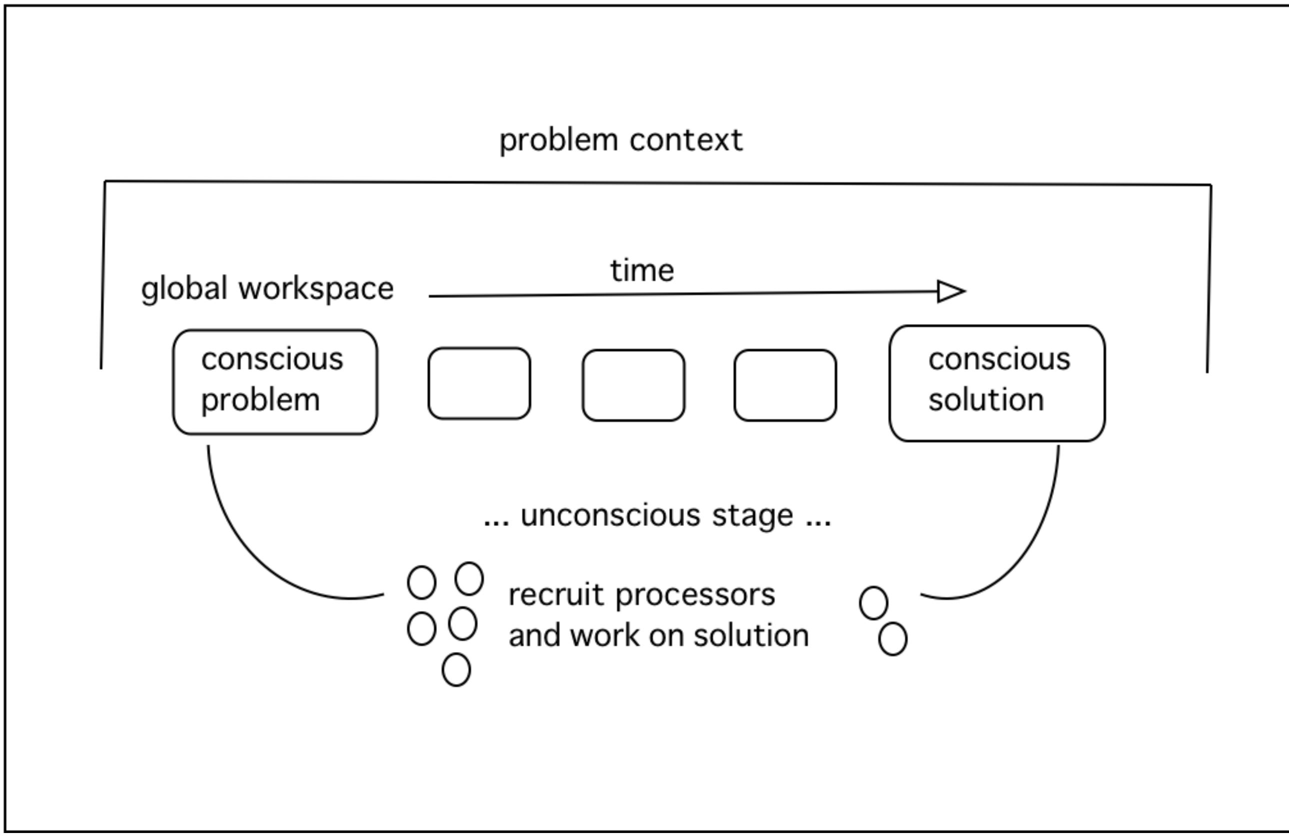

Figure 1. The basic C1-U-C2 triplet of spontaneous problem solving. When a sub-goal is satisfied, the flow of processing can 'pop the stack' to return to higher-level goals. The "tip of the tongue" experience is the simplest example.

Topical threads in the flow of thought.

We can think of each goal-driven string of C1-U-C2 triplets as a *thread* in a set of active, intertwining threads, much like a conversation in an internet chat room. Only a small part of each thread is articulated "in public" at any single moment. But in private, multiple threads are always busy trying to reach their goals. Like consciousness, the "internet chat room" of the mind has quite limited capacity at any moment, and each topical thread appears and disappears in a seemingly arbitrary fashion. The apparent randomness of the free flow of thoughts is of course a common observation in psychology, going back at least to Sir Francis Galton and Sigmund Freud; no doubt an historian could trace it back to the time of Aristotle. (Crovitz, 1970)

The arbitrariness of the stream of thought is of course a central claim in Asian views of consciousness, with ancient roots expressed in Vedanta Hinduism, Buddhism and Taoism. There, the notion of the impulsive or ungovernable mind, which is to be transcended by way of meditation and other practices, is a basic assumption for a sophisticated worldview.

In more recent times the notion of the jumpy randomness of the spontaneous stream of thought has faded in formal science, in good part because naturalistic observation of the stream of thought is methodologically difficult. It is not easy to subject spontaneous thought to experimental study. And yet, unforced thinking continues to exert a fascination. The recent discovery of endogenous “intrinsic networks” of the brain, which are only observed when a subject is freed from the usual experimental demands, is a good example. (Fox et al, 2009).

Our working hypothesis therefore claims that the apparent arbitrariness of the stream of spontaneous thought reflects an underlying pattern; that our spontaneous thinking is guided by a multiplicity of implicit goals and challenges, driven by the most important events in life, our motivations and emotions.

As Jerome Singer has written, mental life is a continuous effort at tracking sensory inputs, cognitively organizing experiences, re-examining memories, and monitoring ‘‘a continuous set of plans and anticipations and a variety of unfinished businesses which compete for our limited attentional capacities with the demands of steering our selves through a physical and social world.’’ (Singer, 1978)

Given this framework, where does spontaneous remembering come in?

2. Hypothesis: Spontaneous remembering is one kind of C-U-C problem solving. Spontaneous thought often seems to have many C-U-C “threads” running at the same time, much like an internet chat discussion. But rather than involving a small group of people, each making their separate contributions, the chat room of consciousness seems to have several major “current concerns” that are touched on in an intertwined way.

Take the case of highly predictable word associations like the following.

1. “brother’; (What is the first association?)

2. “father”: (What is the first association?)

3. “up”: (What is the first association?)’’

Notice that the associate of each word seems to come to mind spontaneously, though it is always highly constrained by lexical, syntactic and semantic regularities. The same point applies to far more complex but still predictable analogies, sentences, concepts, jokes, linguistic strings, musical phrases, and the like.

Remote Associates:

The Remote Associates Test first devised by Mednick provides many good examples. (Mednick, 1962) (http://socrates.berkeley.edu/~kihlstrm/RATest.htm)

Instruction: Think of the first word to come to mind in the blank spaces below.

1. Shopping Washer Picture __________

2. Blank White Lines __________

3. Stick Light Birthday __________

4. Sore Shoulder Sweat ____________

While there are no correct answers for these items, there are very probable and often subjectively surprising associates that come to mind with a strong sense of rightness[1].

The Mednick remote associates have been developed further in recent years to permit brain imaging over many trials. Bowden and colleagues (2005) report a

[1] Reliably high associates to word series are WINDOW, PAPER, CANDLE and COLD respectively.

remarkable interaction of alpha and gamma rhythms during the unconscious incubation period (U, above), followed by the moment of conscious insight (C2, above).

Unconscious semantic inferences evoked by conscious input.

Information can easily be constructed to suggest false inferences, which are largely if not entirely unconscious. Thus,

1. The web consists of …
2. A key board is …

Again there are some surprising answers, that is, answers that appear to violate the spontaneous unconscious inferences we make from these sentence fragments.[2] The effect is similar to the famous "garden path" sentences, which strongly suggest one syntactical structure, only to switch to a different one in the middle of the sentence. A sizable body of such evidence suggests that we spontaneously make unconscious inferences from a wide range of conscious events. In everyday speech, jokes, puns, analogies and insinuations exploit this tendency.

3. Spontaneous problem solving is shaped by implicit motivations and emotions.

What about spontaneous thought? Do we have evidence for similar C1-U-C2

[2] 1. Tiny threads woven by a spider; 2. A wood board to hang keys from.

phenomena there? We propose that the spontaneous stream of thought, and especially its memory component, runs very much along these lines. That is, we propose that spontaneous thought involves conscious moments that trigger unconscious inferences and problem-solving processes, which are then followed by a conscious re-emergence of related material that will tend to complete the C1 and U stages of thought.

Furthermore, we have proposed that many, if not most, spontaneous C-U-C problem-solving in the stream of consciousness is guided by implicit goals and expectations ("goal contexts" in the vocabulary of Baars, 1988). The social psychological evidence for such a claim is now quite strong. (e.g., Fischbach & Shah, 2006; Glaser & Knowles, 2008; Schultheiss et al, 2008).

From psychological evidence to an explicit, large-scale model of cognition.

The LIDA/GWT model generates hypotheses about human cognition by way of its design, the mechanisms of its modules, their interaction, and its performance. All of these hypotheses are, in principle, testable. With the advent of more sophisticated brain and behavioral assessment methods, some earlier hypotheses in this research program have been confirmed (Baars, 2002). We expect the current set of hypotheses to become directly testable with continuing improvements in cognitive neuroscience.

Every autonomous agent (Franklin and Graesser 1997), be it human, animal or artificial, must sample its world and act on it through a sense-select-act (or stimulus, cognition, response) cycle. The LIDA/GWT model hypothesizes for us humans a complex *cognitive cycle*, involving perception, several memory systems, attention and action selection, that samples the world at five to ten times a second. This frequent sampling allows for an exceptionally fine-grained analysis of common cognitive phenomena including spontaneous remembering. At a high level of abstraction, these analyses support the commonly held explanations of what is generally found in studies of the explicit (i.e., conscious and reportable) components of memory processes (e.g. Tulving, 1985; Baddeley et al, 2001). Nothing new here. At a finer grained level, however, our analysis fleshes out these common explanations, adding detail and functional mechanisms. Therein lies the value of our analysis.

In addition, this chapter uses the word “consciousness” or “conscious cognition” to indicate a general cognitive function, much as the word “memory” has come to be used. Conscious cognition is often labeled in many different ways in the empirical literature, including “explicit cognition,” “focal attention,” “awareness,” “strategic processing,” and the like. Here we group all these specific terms under the umbrella of “conscious cognition,” as assessed by standard methods such as verifiable verbal report (Baars 1988). Global Workspace Theory proposes a single underlying kind of information processing for conscious events, as implemented in the LIDA model.

Current techniques for studying these phenomena at a fine grained level, such as PET, fMRI, EEG, implanted electrodes, etc., are still lacking either in scope, in spatial resolution, or in temporal resolution. PET and fMRI have temporal resolution problems, EEG is well-known to have localizability difficulties, and implanted electrodes (in epileptic patients), while excellent in temporal and spatial resolution, can only sample a limited number of neurons; that is, they are limited in scope. As a result, many of our hypotheses, while testable in principle, seem difficult to test at the present time. Improved recording methods are emerging rapidly in cognitive neuroscience (Sigman et al 2007). When GW Theory was first proposed, the core hypothesis of "global activation" or "global broadcasting" was not directly testable in human subjects. Since that time, however, with the advent of brain imaging, widespread brain activation due to conscious, but not unconscious, processes has been found in dozens of studies (see Baars, 2002; Dehaene, 2001). We expect further improvements to make our current hypotheses testable as well.

The LIDA/GWT model has unusual breadth, encompassing perception, working memory, declarative memory, attention, decision-making, procedural learning and more. The model suggests that superficially different aspects of human cognition are so highly integrated that they can't be fully understood in a fragmentary manner. A more global view may provide an overview with surprising points of simplification when analyzing the cognitive mechanisms of spontaneous memory retrieval.

Conscious cognition and memory: Basic facts to be

accounted for

Human memory seems to come in myriad forms: sensory, procedural, working, declarative, episodic, semantic, long-term memory, long-term working memory and many others. How to make sense of all of this? And to add to the difficulty, these terms are used differently in different research traditions. Psychologists tend to use these terms to refer inferentially to systems that appear to hold memory traces and to the underlying knowledge that constitutes their contents. To computer scientists and to neuroscientists, memory refers only to the physical (not inferred) storage device. Further, in many cognitive studies, consciousness is either taken for granted or labeled with its own set of synonyms such as *explicit cognition*, *focal attention*, and *awareness*. Yet the role of consciousness has concerned memory researchers since Ebbinghaus (1885/1964).

There is considerable evidence that people are conscious of retrieved memories in recall, but not necessarily in recognition tasks (e.g., Gardiner et al, 1998). For pioneering memory researchers like Ebbinghaus, indeed, the term “recall” meant retrieval to consciousness. The feeling of knowing that characterizes recognition is a “fringe conscious” phenomenon, that is, an event that has high accuracy but low reported conscious content (Mangan, 2001; Baars, 2002). In numerous experiments, these differences result in striking dissociations between subjective reports in “remember” vs. “know” -types of retrieval.

In cognitive working memory, the active operations of input, rehearsal, recall and report are conscious (Baddeley 1993). The contents of working memory prior to retrieval are not. Baars & Franklin describe the way IDA, an earlier, but compatible version of LIDA/GWT, accounts for this evidence. (2003).

Novel Hypotheses from the LIDA/GWT Model

With its finer-grained model of these processes, the LIDA model (Franklin 2000, 2001; Franklin & Graesser 2001; Ramamurthy, D'Mello and Franklin. 2006) of Global Workspace Theory (Baars 1988, 1993, 1997, 2002) offers hypotheses that suggest a simple account of several forms of human memory and their relationships with conscious events, including spontaneous memories. Here we list, and briefly discuss, several of these relevant hypotheses.

1. **The Cognitive Cycle**: Recall William James' claim that the stream of conscious thought consists of momentary "flights" and somewhat longer "perches" of dwelling on a particular conscious event. Such findings have been reported by neuroscientists (Halgren et al 2002, Fuster et al, 2000; Lehmann et al 1998; Freeman, 2003). Much of human cognition functions by means of continual interactions between conscious contents, the various memory systems and decision-making. We call these interactions, as modeled in LIDA, *cognitive cycles.* While these cycles can overlap, producing cascading processes, they must preserve the seriality of consciousness. The LIDA model suggests therefore that *conscious events occur as a sequence of discrete, coherent episodes separated by quite*

short periods of no conscious content[3] (see also VanRullen and Koch 2003). It should be pointed out that the "flights and perches" of normal consciousness may involve numerous cognitive cycles. A problem-solving task involving inner speech, for example, may occur over tens of seconds or minutes, according to careful thought-monitoring studies.

2. **Transient Episodic Memory**: Humans have a content-addressable, associative, transient episodic memory with a decay rate measured in hours (Conway 2001). In our theory, *a conscious event is stored in transient episodic memory by a broadcast from a global workspace*. A corollary to this hypothesis says that *conscious contents can only be encoded (consolidated) in long-term declarative memory via transient episodic memory.*

3. **Perceptual Memory**: *A perceptual memory, distinct from semantic memory but storing some of the same contents, exists in humans, and plays a central role in the assigning of interpretations to incoming stimuli.* The conscious broadcast begins and updates the process of learning to recognize and to categorize, both employing perceptual memory.

---

3 To say that conscious moments may be separated by short periods of no conscious contents does not mean, of course, that people momentarily fall asleep between conscious episodes.

4. **Consciousness**: *Conscious cognition is implemented computationally by way of a broadcast of contents from a global workspace*, which receives input from the senses and from memory (Baars 2002).

5. **Conscious Learning**: *Significant learning takes place via the interaction of consciousness with the various memory systems* (e.g. Standing 1973; Baddeley, 1993). The effect size of subliminal learning is therefore small compared to the learning of conscious events, but significant implicit learning can occur by way of unconscious inferences based on conscious patterns of input (Reber et al, 1991). In the LIDA/GWT view, all memory systems represented in the model rely on conscious cognition for their updating, either in the course of a single cycle or over multiple cycles.

6. **Voluntary and Automatic Attention:** In the LIDA/GWT model, attention is defined as the process of bringing contents to consciousness. *Automatic attention may occur unconsciously and without effort, even during a single cognitive cycle* (Logan 1992). Attention may also occur voluntarily and *effortfully* in a conscious, goal-directed way, over multiple cycles.

7. **Voluntary and Automatic Memory Retrievals:** *Associations from transient episodic and declarative memory are retrieved automatically and unconsciously during each cognitive cycle.* Voluntary retrieval from these memory systems may occur over multiple cycles, governed by conscious goals.

Global Workspace Theory as a Functional Interpretation of Conscious Cognition.

Global workspace theory is a cognitive architecture with an explicit role for consciousness. It makes the following assumptions:

1. The brain may be viewed as a collection of distributed *specialized networks* (processors);

2. Consciousness is associated with a *global workspace* -- a fleeting memory capacity whose focal contents are widely distributed ("broadcast") to many unconscious specialized networks;

3. Some unconscious networks, called *contexts*, shape conscious contents (for example, unconscious parietal maps of the visual field modulate feature cells needed for conscious vision);

4. Such contexts may work together to jointly constrain conscious events;

5. Motives, implemented by feelings and emotions[4], can be viewed as part of *goal contexts, which are often unconscious*;

6. Voluntary control employs *hierarchies of goal contexts*.

---

4 Feelings in humans include hunger, thirst, various sorts of pain, hot or cold, the urge to urinate, tiredness, depression, etc. Emotions, such as fear, anger, joy, sadness, shame, embarrassment, resentment, regret, guilt, etc., are taken to be feelings with cognitive content (Johnston 1999).

A number of these functions have plausible brain correlates, and the theory has recently gathered considerable interest from cognitive neuroscience and philosophy (Cooney & Gazzaniga 2003, Damasio 1989, Dehaene & Naccache 2001, Edelman & Tononi 2000, Freeman 2003a, Varela et al 2001)[5]. For instance, Dennett notes that "Theorists are converging from quite different quarters on a version of the global neuronal workspace model of consciousness…" (2001).

The LIDA Model

The LIDA model is a comprehensive, conceptual and computational[6] model covering a large portion of human cognition. Based primarily on global workspace theory (Baars 1988, 1997, 2002), the model implements and fleshes out a number of psychological and neuropsychological theories including situated cognition (Varela et al. 1991), perceptual symbol systems (Barsalou 1999), working memory (Baddeley and Hitch 1974, Baddeley 1993), memory by affordances[7] (Glenberg 1997), long-term working memory (Ericsson and Kintsch 1995), transient episodic memory (Conway 2001), and the H-CogAff framework (Sloman 1999). The LIDA model is particularly compatible with the notion of grounded cognition (Barsalou 2008). LIDA's flexible cognitive cycle has been

5 Quotes to this effect from each of these citations, and more, can be found in Baars, 2002 and in Franklin et al, 2005.

6 At this writing the LIDA model is only partially implemented. We claim it as a computational model since each of its modules and most of its processes have been designed for implementation.

7 Gibson (1979) introduced the term *affordance*, meaning that information about the available uses of an object existed in the object itself. We are using it in the sense that the agent can derive such information from the object.

used to analyze the relationship of consciousness to working memory at a fine level of detail, offering explanations of such classical working memory tasks as the “phonological loop” to account for the rehearsal of a telephone number (Baars and Franklin 2003). There is evidence suggesting such a cognitive cycle from neurobiology in the form of “…hemisphere-wide, self-organized patterns of perceptual neural activity…” recurring aperiodically at intervals of 100 to 200 ms (Freeman 2003, Freeman 2003a, Lehmann et al 1998, Koenig, Kochi and Lehmann 1998).

Memory Systems and Terminology

In this section, we will briefly discuss the various human memory systems that will play a role in the rest of the article. It will be helpful to the reader to specify here how we plan to use the various terms, as there isn’t always agreement in the literature. Figure 2 displays some of the relationships between the memory systems we’ll discuss.

Sensory memory holds incoming sensory data in sensory registers and is relatively unprocessed. In addition to deriving a representation of geometric properties of the current situation, it provides a workspace for integrating the features from which representations of objects and their relations are constructed. It also sends information along the dorsal stream to facilitate the executing of actions. There are different sensory memory registers for different senses: *iconic* (visual), *echoic*, *haptic*, and likely a separate *sensory memory* for integrating multimodal information. Sensory memory has the fastest decay rate, measured in tens of milliseconds.

Working memory is the manipulable scratchpad of the mind (Miyake and Shah 1999). It holds sensory data, both endogenous (for example, visual images and inner speech) and exogenous (sensory), together with their interpretations. Its decay rate is measured in tens of seconds. Again, there are separate working memory components associated with the different senses, the visuo-spatial sketchpad and the *phonological loop* for example (Baddeley 1993; Franklin & Baars 2003). Also, there are long-term processing components of working memory (Ericsson and Kintsch 1995). Baars & Franklin (2003) have suggested that conscious input, rehearsal, and retrieval are necessary for the normal functions of working memory.

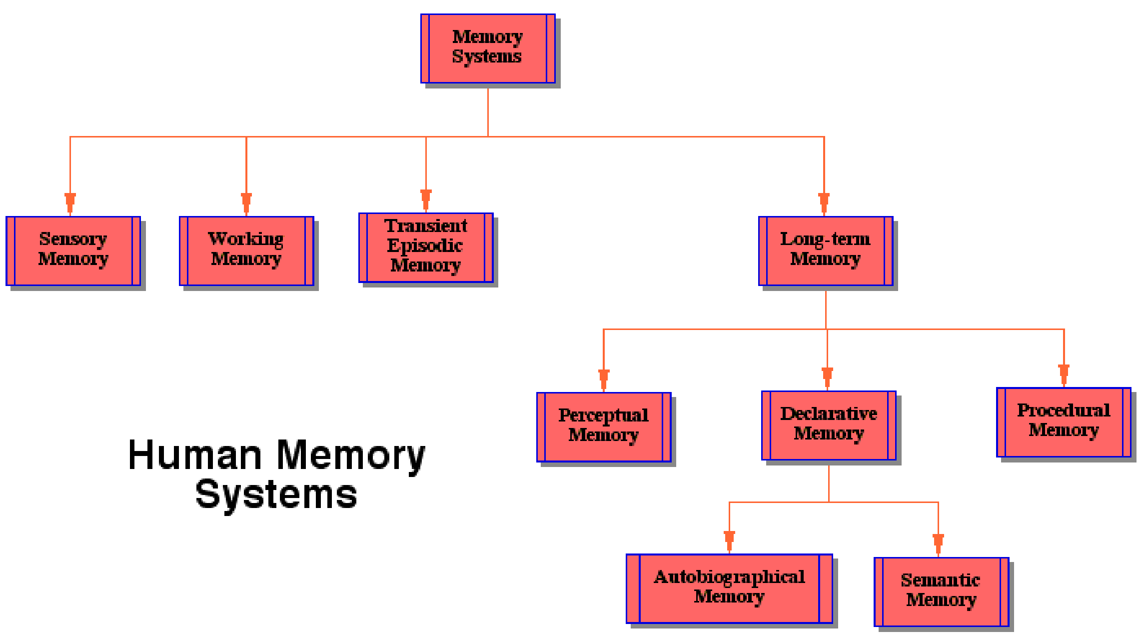

Figure 2. Human Memory Systems

Episodic or autobiographical memory is memory for events having features of a particular time and place (Baddeley, Conway and Aggleton 2001). This memory system is associative and content-addressable.

An unusual aspect of the LIDA model is transient episodic memory (TEM), an episodic memory with a decay rate measured in hours. Though often assumed (Panksepp 1998, p 129, assumes a “transient memory store”), the existence of such a memory has rarely been explicitly asserted (Donald 2001; Conway 2001; Baars and Franklin 2003). It will play a major role in the hypotheses about memory systems generated by the LIDA model.

Humans are blessed with a variety of long-term memory types that can decay exceedingly slowly, if at all. Memory researchers typically distinguish between procedural memory, the memory for motor skills including verbal skills, and declarative memory. In the LIDA model, declarative memory (DM) is composed of autobiographical memory, described in a previous paragraph, and semantic memory, memories of fact or belief typically lacking a particular source with a time and place of acquisition. In contrast, semantic memories have lost their association with their original autobiographical source. DM is a single system within the LIDA model. These declarative memory systems are accessed by means of specific cues from working memory. The LIDA model hypothesizes that DM decays inversely with the strength of the memory traces.

Though “perceptual memory” is often used synonymously with “sensory memory,” we follow Taylor (1999 p. 29) and use the term differently (see also

Barsalou 2008). Perceptual memory is a memory for individuals, categories, and their relations. The LIDA model distinguishes between semantic memory and perceptual memory (PM) and hypothesizes distinct mechanisms for each (Nadel 1992, Franklin et al. 2005). According to the model, PM plays the major role in recognition, categorization, and more generally the assignment of interpretations, for example the recognition of situations. Upon presentation of features of an incoming stimulus, PM returns interpretations. The content of semantic memory is hypothesized to be a superset of that of PM. All this discussion essentially restates the most controversial part of our *Perceptual Memory Hypothesis*, the claim of distinct mechanisms for PM and semantic memory. Several types of evidence, of varying degrees of persuasiveness, support this dissociation (Franklin et al 2005), including arguments from evolution, from developmental studies, from clinical studies of amnesiacs, and from experiments with animals with their hippocampal systems excised.

In the recognition memory literature dual-process models have been put forward proposing that two distinct memory processes, referred to as *familiarity* and *recollection*, support recognition (Mandler 1980, Jacoby and Dallas 1981). Familiarity allows one to recognize the butcher in the subway *acontextually* as someone who is known, but not to recollect the context of the butcher shop. In the LIDA model, PM alone provides the mechanism for such a familiarity judgment, while both PM and DM are typically required for recollection. Recent brain imaging results from cognitive neuroscience support a dual-process model (Rugg and Yonelinas 2003), and so are compatible with our Perceptual Memory

Hypotheses.

The LIDA Cognitive Cycle

The LIDA model and its ensuing architecture are grounded in the LIDA cognitive cycle. Every autonomous agent (Franklin and Graesser 1997), be it human, animal, or artificial, must frequently sample (sense) its environment and select an appropriate response (action). More sophisticated agents, such as humans, process (make sense of) the input from such sampling in order to facilitate their decision-making. Neuroscientists call this three-part process the action-perception cycle (Freeman 2002). The agent's "life" can be viewed as consisting of a continual cascading sequence of these cognitive cycles. Each cycle constitutes a unit of sensing, attending and acting.

A cognitive cycle can be thought of as a moment of cognition -- a cognitive "moment." Higher-level cognitive processes are composed of many of these cognitive cycles, each a cognitive "atom." Citing evidence from Thompson *et al.* (1996) and from Skarda and Freeman (1987), Cotterill speaks of "… the time usually envisioned for an elementary cognitive event … about 200 ms" (2003). From our cognitive cycle hypothesis, it might seem reasonable to call one such cycle an elementary cognitive event. Freeman (1999) suggests that conscious events succeed one another at a "frame rate" of 6 Hz to 10 Hz as would be expected from our cognitive cycle hypothesis (See also Freeman 2003b). The rate of such cycles coincides roughly with that of other, perhaps related, biological cycles such as *saccades* (Steinman, Kowler and Collewijn 1990), *systematic* *motor tremors*[8], and *vocal vibrato* (Seashore 1967). Could these hypothesized cycles be related to hippocampal theta waves (at 6-9 hz) with gamma activity superimposed on them (VanRullen and Koch. 2003)?

Just as atoms are composed of protons, neutrons and electrons, and some of these are composed of quarks, bosons, muons, etc., these cognitive "atoms" have a rich inner structure. What the LIDA model hypothesizes as this rich inner structure of the LIDA cognitive cycle will now be described. More detailed descriptions are available elsewhere (Baars and Franklin 2003, Franklin et al. 2005).

During each cognitive cycle the LIDA agent first makes sense of its current situation as best as it can *by updating its representation of its current situational model, both external and internal*. By a competitive process to be described below, it then decides what portion of the represented situation is most in need of attention. Broadcasting this portion, the current contents of consciousness, enables the agent to finally chose an appropriate action and execute it. Let's look at these three processes in a little more detail. Figure 3 should help the reader follow the description. It starts in the upper left corner and proceeds roughly clockwise.

8 "In the last 15 years or so, it has become clear that the 8-12Hz rhythmicity of physiological tremor is observed not only during voluntary movement, but also … during maintained posture and in supported limbs at rest (Marsden et al. 1984)." (Llinas 2001)

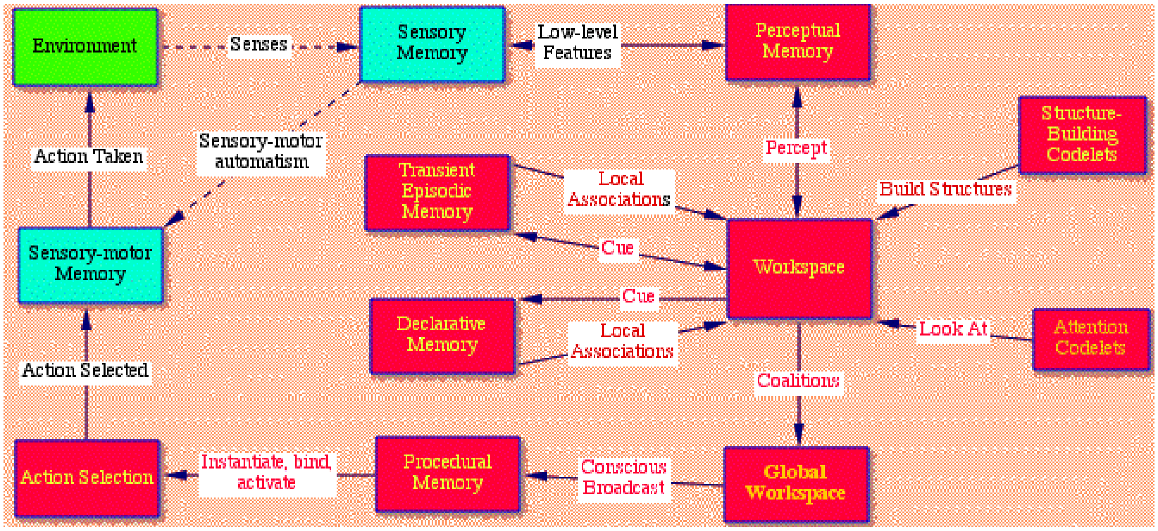

Figure 3: LIDA Cognitive Cycle Diagram

The cycle begins with sensory stimuli from the agent's environment, both an external and an internal environment. Low-level feature detectors in sensory memory begin the process of making sense of the incoming stimuli. These low-level features are passed on to perceptual memory where higher-level features, such as objects, categories, relations, situations, etc. are recognized. These entities, which have been recognized preconsciously, make up the percept that is passed to the workspace, where a model of the agent's current situation is assembled. (LIDA's workspace contains the preconscious buffers of working memory.) This percept serves as a cue to two forms of episodic memory, transient and declarative. Responses to the cue consist of local associations, that is, remembered events from these two memory systems that were associated with the various elements of the cue. In addition to the current percept, the workspace contains recent percepts and the models assembled from them that haven't yet decayed away.

A new model of the agent's current situation is assembled from the percepts, the associations, and the undecayed parts of the previous model. This assembling process will typically require structure-building codelets[9]. These structure-building codelets are small, special purpose processors, each of which has some particular type of structure it is designed to build. To fulfill their task these codelets may draw upon perceptual memory and even sensory memory, to enable the recognition of relations and situations. The newly assembled model constitutes the agent's understanding of its current situation within its world. It has made sense of the incoming stimuli.

For an agent "living" in a complex, dynamically changing environment, this current model may well be much too much for the agent to consider all at once in deciding what to do next. It needs to select what portion of the model should be attended to. Which are the most relevant, important, urgent or insistent structures within the model? Portions of the model compete for attention. These competing portions take the form of coalitions of structures from the model. Such coalitions are formed by attention codelets, whose function is to bring certain structures to consciousness. One of the coalitions wins the competition. In effect, the agent has decided on what to attend.

However, the purpose of all this processing is to help the agent decide what to do next. To this end, a representation of the contents of the winning coalition is

9 The term codelet refers generally to any small, special purpose processor or running piece of computer code. The concept is essentially the same as Baars' processors (1988), Minsky's agents (1985), Jackson's demons (1987), or Ornstein's small minds (1986). The term was borrowed from Hofstadter and Mitchell (1995).

broadcast globally, effectively constituting a global workspace, (hence the name global workspace theory). Though the contents of this conscious broadcast are available globally, the primary recipient is procedural memory, which stores templates of possible actions including their contexts and possible results. It also stores an activation value for each such template that attempts to measure the likelihood of an action taken within its context producing the expected result. Templates whose contexts intersect sufficiently with the contents of the conscious broadcast instantiate copies of themselves with their variables specified to the current situation. Instantiated templates remaining from previous cycles may also continue to be available. These instantiations are passed to the action selection mechanism, which chooses a single action from one of these instantiations. The chosen action then goes to sensory-motor memory, where it picks up the appropriate algorithm by which it is then executed. The action taken affects the environment, and the cycle is complete.

The LIDA model hypothesizes that all human cognitive processing is via a continuing iteration of such cognitive cycles. These cycles occur asynchronously, with each cognitive cycle taking roughly 200 ms. The cycles cascade, that is, several cycles may have different processes running simultaneously in parallel. This cascading must, however respect the way consciousness processes information serially in order to maintain the stable, coherent image of the world with which consciousness endows us (Merker 2005, Franklin 2005). This cascading, together with the asynchrony, allows a rate of cycling in humans of five to ten cycles per second. A cognitive “moment” is thus quite short! There is considerable empirical evidence from neuroscience suggestive of and consistent with such cognitive cycling in humans (Massimini et al. 2005, Sigman and Dehaene 2006, Uchida, Kepecs and Mainen 2006, Willis and Todorov 2006). None of this evidence is conclusive, however.

Learning in the LIDA model

Edelman (1987) usefully distinguishes two forms of learning, the selectionist and the instructionalist. Selectionist learning requires selection from a redundant repertoire that is typically organized by some form of reinforcement learning. A repertoire of, say, possible actions, is redundant if slightly different actions can lead to roughly the same result. In reinforcement learning (Kaelbling, Littman and Moore 1996) a successfully executed action is reinforced, making it more likely to be chosen the next time the result in question is needed. In Edelman’s system little used actions tend to decay away. In contrast, instructional learning allows the learning of representations of, say, new actions, that is, actions not currently in the repertoire.

Global workspace theory postulates that learning requires only attention (Baars 1988 pp. 213-218). In the LIDA model this implies that learning must occur with each cognitive cycle, because whatever enters consciousness is being attended to. More specifically, learning occurs with the conscious broadcast from the global workspace during each cycle. Learning in the LIDA model follows the tried and true Artificial Intelligence principle of generate and test. New representations are learned in a profligate manner (the generation) during each cognitive cycle. Those that are not sufficiently reinforced during subsequent cycles (the test) decay away. Three modes of learning -- perceptual, episodic and procedural -- employing distinct mechanisms (Nadel 1992, Franklin et al. 2005), have been designed and are in various stages of implementation. A fourth, attentional learning, is contemplated but not yet designed. We'll discuss each individually.

Perceptual learning enables an agent to recognize features, objects, categories relations, situations, etc. In the LIDA model what is learned perceptually is stored in perceptual memory (Franklin 2005a, Franklin 2005b). Motivated by the Slipnet from the Copycat architecture (Hofstadter and Mitchell 1995), the LIDA perceptual memory is implemented as a collection of nodes and links with activation passing between the nodes. Nodes represent features, individuals, categories, actions, feelings and more complex structures. Links, both excitatory and inhibitory, represent relations. Each node and link has both a current and a base-level activation. The base-level activation measures how useful the node or link has been in the past, while the current activation depends on its relevance in the current situation. The percept passed on to the workspace during each cognitive cycle is composed of those nodes and links whose total activation is over the threshold. Perceptual learning in its selectionist form modifies base-level activation, and in its instructionalist form creates new nodes and links. One or the other or both may occur with the conscious broadcast during each cognitive cycle.

Episodic learning refers to the memory of events -- the what, the where and the when (Tulving 1983, Baddeley, Conway and Aggleton 2001). In the LIDA model such learned events are stored in transient episodic memory (Conway 2001, Franklin et al. 2005) and in the longer-term declarative memory (Franklin et al. 2005). Both are implemented using sparse distributed memory (Kanerva 1988), which is both associative and content addressable and has other desirable psychological properties. In particular it knows when it doesn't know, and exhibits the tip of the tongue phenomenon. Episodic learning in the LIDA model (Ramamurthy, D'Mello and Franklin 2004, Ramamurthy, D'Mello and Franklin 2005) is also a matter of generate and test, with such learning occurring at the conscious broadcast of each cognitive cycle. Episodic learning is initially directed only to transient episodic memory. At a later time and offline, the undecayed contents of transient episodic memory are consolidated (Nadel and Moscovitch. 1997, Stickgold and Walker 2005) into declarative memory, where they still may decay away or may last a lifetime.

Procedural learning refers to the learning of new tasks and the improvement of old tasks. In the LIDA model such learning is accomplished in procedural memory (D'Mello et al. 2006), which is implemented via a scheme net motivated by Drescher's schema mechanism (Drescher, 1991). Each scheme in procedural memory is a template for an action, consisting of a context, an action and a result, together with a base-level activation intended to measure how likely the result would be to occur were the action taken within its specific context. Once again, the LIDA model's procedural learning is via a generate and test mechanism, using base-level activation as reinforcement, as well as through the creation of new schemes. These new schemes can support multiple actions, both parallel and sequential.

The Availability Heuristic

It is well known that people tend to overestimate the frequency of divorce if they can quickly recall instances of divorced acquaintances. This principle also applies to frequency estimates, and is referred to as the *availability heuristic* (Kahneman, Slovic and Tversky 1982; Fiske and Taylor 1991).

An online demonstration of the heuristic (Colston and Walter. 2001) asks the subject to review a list of names of well-known people, one such presented at each mouse click, to see if the subject knows them. No mention is made of gender. After viewing the last name in the list, the subject is presented a forced choice as to whether he or she had seen more men’s names or women’s. Since the men named tended to be more famous, and hence more easily recalled, the availability heuristic would correctly predict that most subjects would claim that there were more men’s names on the list. There are, in fact, fourteen of each. In this section, we will analyze this task using LIDA’s cognitive cycle to see what the LIDA model would predict for a human subject.

The initial instructions given to the subject comprise a text of some thirty-seven words. To read and understand the instructions will likely occupy a subject for a few seconds and some few tens of cognitive cycles. During the last of these cycles, the gist of the meaning of the instructions will have accumulated in the appropriate preconscious working memory buffer. The conscious broadcast of these meanings will likely instantiate a goal context hierarchy for sequentially clicking through and seeing the names, noting whether they are recognized. The action chosen during this cycle will likely be a mouse click bringing up the first name.

One to three or four cycles will likely suffice for the subject to preconsciously perceive the entire name, which will have accumulated in the appropriate preconscious working memory buffer in LIDA's workspace. Here we must consider two cases: 1) The name has been recognized during the preconscious perceptual process (see the description of perceptual memory above), or 2) it has not. Recognition of the name John Doe has occurred if the subject can answer the question "who is John Doe?" In the first case, after the conscious broadcast the action selection mechanism will likely choose a mouse click as an action. This would bring up another name. In the second case, a conscious goal to consult declarative memory in search of recognition would likely arise over several cognitive cycles (see major novel hypothesis 7 on voluntary memory retrieval above.) Such a conscious goal would produce an attention codelet on the lookout for information concerning the as yet unrecognized name. In the next cycle, the name, by then located in a preconscious working memory buffer, will be used to cue declarative memory. This voluntary episodic memory retrieval process may iterate over several additional cycles, with the parts of local associations that make it to consciousness contributing to the cue for the next local associations. The subject will eventually recognize the name or will give up the effort on a subsequent cycle. In either case, the action chosen selected on this last cycle is likely to be a mouse click for the next name.

Thus the subject will work his or her way through the list of names, recognizing many or most of them but missing some. At the mouse click following the last name, a new set of instructions appears. These instructions, in sixteen words, ask the subject to decide whether more men's or women's names were on the list, and to click the mouse when the decision is made. A very few tens of cycles are spent understanding the instructions.

The conscious broadcasts whose contents contain the full understanding of the instructions will likely recruit behavior codelets that instantiate a goal context hierarchy to comply. During some subsequent cycles, behaviors (goal contexts) will likely be selected that will attempt to query transient episodic memory (TEM), starting the recall process for names on the list. The behavior's codelets will write this goal to preconscious buffers of working memory (WM) where it will serve to cue local associations from TEM and declarative memory (DM). The next contents of consciousness will be chosen from the resulting long-term working memory (LTWM), a part of LIDA's workspace.

This process can be expected to continue over many subsequent cycles, with each cue from WM containing material from previous local associations from DM. Since the names were encoded in TEM as distinct events, those retrieved from TEM can be expected to appear as such in LTWM when associated with the latest cue. From there, each recovered name is likely to appear as the central content of its own coalition of that some attention and information codelet brings tos competing compete for consciousness. Thus, the LIDA model would predict that only one name would be recalled at a time, since a single coalition must win the competition. More famous names would have their initial coalitions replaced during subsequent cycles, with expanded coalitions including information from their many local associations in DM. Thus, they would accrue an advantage in the next competition for consciousness; such, so that more famous names are more likely to be recalled. A name on the list could fail to be recalled for either of two reasons. One, it might not have been retrieved from TEM by any of the cues used during the process. Or, it may have been retrieved into LTWM, but decayed away before it could be part of a coalition that won the competition for consciousness.

At some point in this process, a decision as to more men's names or more women's is taken. How does that happen? Some subjects may actually instantiate a goal context hierarchy (behavior stream) that keeps separate tallies in WM of the number of men's names and of women's names recalled. These tallies are likely to be part of the conscious contents in at least some of the ongoing cycles in the process. When no more names are being recalled, a goal context (behavior) to decide is chosen. The decision is then made and, on a subsequent cycle, the mouse is clicked.

We think this is a relatively unlikely scenario, in that most subjects will not keep such explicit running tallies. Rather, they'll decide on the basis of a fringe consciousness feeling that one gender or the other has been recalled more often

(Mangan 2001). In the LIDA model, such feelings are to be implemented as fringe attention codelets. In this case, a selected behavior would give rise to two such fringe attention codelets, one for each gender. As a name from one gender is recalled, the activation of the corresponding fringe attention codelet is increased, the amount of increase biased by the activation of the recall of the name. Each of these fringe attention codelets likely enters each competition for consciousness as the process progresses. The stronger of the two will be able to win only after names are no longer being recalled; a coalition with a name to be newly recalled would simply have a higher average activation. Such fringe consciousness feelings are easily defeated.

Thus the LIDA model predicts the experienced outcome of this demonstration of the availability heuristic, and supports the commonly given, functional explanation, the availability heuristic. *What the model adds is a hypothesized detailed mechanism for its functional process.*

Conclusions

In the above we have explore some of the well-established features of spontaneous remembering in the light of the LIDA/GWT model, including incidental learning and problem solving, expectation-driven recall, action control, and the "availability heuristic." Our basic message is that spontaneous remembering is a part of our ordinary stream of thought, helping us to answer the ever-present question: *What do I do next, to further my most important goals?* This view suggests that spontaneous thoughts are not just irrelevant "mind-wandering," as some researchers suggest, but rather a highly functional, implicitly purposeful, problem-solving stream.

Acknowledgements

The authors gratefully acknowledge the conceptual contributions to this work made by the various members, past and present, of the Cognitive Computing Research Group at the University of Memphis, and to Uma Ramamurthy and Mathew Ventura in particular.

References

Baars, B. J. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Baars, B. J. 1993. How does a serial, integrated and very limited stream of consciousness emerge from a nervous system that is mostly unconscious, distributed, and of enormous capacity? In *CIBA Symposium on Experimental and Theoretical Studies of Consciousness*, ed. G. R. Bock, and J. Marsh. London: John Wiley and Sons.

Baars, B. J. 1997. *In the Theater of Consciousness*. Oxford: Oxford University Press.

Baars, B. J. 2002. The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Science* 6:47–52.

Baars, B. J., and S. Franklin. 2003. How conscious experience and working memory interact. *Trends in Cognitive Science* 7:166-172.

Baddeley, A. D. 1993. Working memory and conscious awareness. In *Theories of memory,* ed. A. Collins, S. Gathercole, M. Conway, and P. Morris. Howe: Erlbaum.

Baddeley, A., M. Conway, and J. Aggleton. 2001. *Episodic Memory*. Oxford: Oxford University Press.

Baddeley, A. D., and G. J. Hitch. 1974. Working memory. In The Psychology of Learning and Motivation, ed. G. A. Bower. New York: Academic Press.

Barsalou, L. W. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22:577–609.

Barsalou, LawrenceW. 2008. Grounded cognition. *Annual Review of Psychology* 59: 617–45.

Bowden, Edward M., Mark Jung-Beeman, Jessica Fleck and John Kounios. 2005. New approaches to demystifying insight, *Trends in Cognitive Sciences*, 9(7), 322-328.

Chafe, Wallace 2000. A linguist's perspective on William James and “the stream of thought.” *Consciousness and Cognition* 9, no. 4: 618-628.

Conway, M. A. 2001. Sensory-perceptual episodic memory and its context: autobiographical memory. In *Episodic Memory*, ed. A. Baddeley, M. Conway, and J. Aggleton. Oxford: Oxford University Press.

Cooney, J. W., and M. S. Gazzaniga. 2003. Neurological disorders and the structure of human consciousness. *Trends in Cognitive Sciences* 7:161-165.

Crovitz, Herbert F. 1970. Galton's Walk. New York; Harper & Row.

Damasio, A. R. 1989. Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33:25.62.

Dehaene, S. 2001. *The Cognitive Neuroscience of Consciousness*. Cambridge, MA; Amsterdam: MIT Press; Elsevier.

Dehaene, S., and L. Naccache. 2001. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 79:1-37.

D'Mello, S. K., U. Ramamurthy, A. Negatu, and S. Franklin. 2006. A Procedural Learning Mechanism for Novel Skill Acquisition. In *Workshop on Motor Development*: *Proceeding of Adaptation in Artificial and Biological Systems, AISB'06*, vol. 1, ed. T. Kovacs, and J. A. R. Marshall. Bristol, England: Society for the Study of Artificial Intelligence and the Simulation of Behaviour; April 2006.

Dennett, D. C. 2001. Are we explaining consciousness yet? *Cognition* 79:221-237.

Donald, M. 2001. *A Mind So Rare*. New York: Norton.Edelman, G. M. 1987. *Neural Darwinism*. New York: Basic Books.

Drescher, G. L. 1991. Made-Up Minds: A Constructivist Approach to Artificial Intelligence. Cambridge, MA: MIT Press.

Ebbinghaus, H. 1885/1964. *Memory: A contribution to experimental psychology*. New York: Dover.

Edelman, G. M. 1987. *Neural Darwinism*. New York: Basic Books.

Edelman, G. M., and G. Tononi. 2000. *A Universe of Consciousness*. New York: Basic Books.

Epstein, Russell 2000. The neural-cognitive basis of the Jamesian stream of thought. *Consciousness and Cognition* 9: 550–575.

Ericsson, K. A., and W. Kintsch. 1995. Long-term working memory. *Psychological Review* 102:21–245.

Fishbach, A., Shah, J. Y. 2006. Self-Control in Action: Implicit Dispositions Toward Goals and Away From Temptations, *Journal of Personality and Social Psychology*, 90(5): 820-832.

Fox MD, Zhang D, Snyder AZ, Raichle ME. 2009. The Global Signal and Observed Anticorrelated Resting State Brain Networks. *J Neurophysiol*. 2009 Apr 1.

Franklin, S. 2000. Modeling Consciousness and Cognition in Software Agents. In *Proceedings of the Third International Conference on Cognitive Modeling, Groeningen, NL, March 2000*, ed. N. Taatgen. Veenendal, NL: Universal Press.

Franklin, S. 2001. Conscious Software: A Computational View of Mind. In *Soft Computing Agents: New Trends for Designing Autonomous Systems*, ed. V. Loia, and S. Sessa. Berlin: Springer (Physica-Verlag).

Franklin, Stan. 2005. Evolutionary pressures and a stable world for animals and robots: A commentary on Merker. *Consciousness and Cognition* 14: 115–118.

Franklin, S. 2005a. Cognitive Robots: Perceptual associative memory and learning. In Proceedings of the 14th Annual International Workshop on Robot and Human Interactive Communication (RO-MAN 2005).

Franklin, S. 2005b. Perceptual Memory and Learning: Recognizing, Categorizing, and Relating. Symposium on Developmental Robotics; American Association for Artificial Intelligence (AAAI); Stanford University, Palo Alto CA, USA. March 21-23, 2005.

Franklin, S, B J Baars, U Ramamurthy, and Matthew Ventura. 2005. The role of consciousness in memory. *Brains, Minds and Media* 1: 1–38.

Franklin, S., and A. C. Graesser. 1997. Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents. In *Intelligent Agents III*. Berlin: Springer Verlag. 21-35.

Franklin, S., and A. Graesser. 2001. Modeling Cognition with Software Agents. In *CogSci2001*: *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, ed. J. D. Moore, and K. Stenning. Mahwah, NJ: Lawrence Erlbaum Associates; August 1-4, 2001.

Franklin, S., A. Kelemen, and L. McCauley. 1998. IDA: A Cognitive Agent Architecture. In *IEEE Conf on Systems, Man and Cybernetics*. : IEEE Press.

Franklin, Stan and F G Jr Patterson. 2006. The LIDA architecture: Adding new modes of learning to an intelligent, autonomous, software agent. In *IDPT-2006 proceedings (integrated design and process technology)*: Society for Design and Process Science.

Freeman, W. J. 1999. Consciousness, Intentionality and Causality. In *Reclaiming Cognition*, ed. R. Núñez, and W. J. Freeman. Exeter: Imprint Academic.

Freeman, W J. 2002. The limbic action-perception cycle controlling goal-directed animal behavior. *Neural Networks* 3: 2249-2254.

Freeman, W. J. 2003. Evidence from human scalp electroencephalograms of global chaotic itinerancy. *Chaos* 13:1067-1077.

Freeman, W. J. 2003a. The wave packet: An action potential for the 21st century. *Journal of Integrative Neuroscience* 2: 3–30.

Freeman, W. J. 2003b. High resolution EEG brings us another step closer to the NCC?. Plenary. ASSC7: Seventh Conference of the Association for the Scientific Study of Consciousness; Association for the Scientific Study of Consciousness; Memphis, TN, USA; May 30th to June 2nd, 2003.

Fuster, J. M., M. Bodner, and J. K. Kroger. 2000. Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature* 405:347-351.

Gardiner, J. M., C. Ramponi, and A. Richardson-Klavehn. 1998. Experiences of remembering, knowing, and guessing. *Consciousness and Cognition* 7:1-26.

Gibson, James J. 1979. *The ecological approach to visual perception*. Mahwah, New Jersey: Lawrence Erlbaum Associates.

Glaser, J., Knowles, E.D. 2008. Implicit motivation to control prejudice, *Journal of Experimental Social Psychology*, 44(1): 164-172

Glenberg, A. M. 1997. What memory is for. *Behavioral and Brain Sciences* 20:1–19.

Gray, J. A. 1995. The contents of consciousness: a neuropsychological conjecture. *Behavioral and Brain Sciences* 18:659–676.

Halgren, E., C. Boujon, J. Clarke, C. Wang, and P. Chauvel. 2002. Rapid distributed fronto-parieto-occipital processing stages during working memory in humans. *Cerebral Cortex* 12:710-728.

Hofstadter, D. R., and M. Mitchell. 1994. The Copycat Project: A model of mental fluidity and analogy-making. In *Advances in connectionist and neural computation theory, Vol. 2: logical connections*, ed. K. J. Holyoak, and J. A. Barnden. Norwood N.J.: Ablex.

Jackson, J. V. 1987. Idea for a Mind. *Siggart* Newsletter, 181:23–26.

James, W. 1890. *The Principles of Psychology*. Cambridge, MA: Harvard University Press.

Johnston, V. S. 1999. *Why We Feel:The Science of Human Emotions*. Reading MA: Perseus Books.

Kaelbling, L. P., M. L. Littman, and A. W. Moore. 1996. Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research* 4:237-285.

Kanerva, P. 1988. *Sparse Distributed Memory*. Cambridge MA: The MIT Press.

Koenig, T., K. Kochi, and D. Lehmann. 1998. Event-related electric microstates of the brain differ between words with visual and abstract meaning. *Electroencephalogr. Clin. Neurophysiol.* 106:535-546.

Lehmann, D., W. K. Strik, B. Henggeler, T. Koenig, and M. Koukkou. 1998. Brain electric microstates and momentary conscious mind states as building blocks of spontaneous thinking: I. Visual imagery and abstract thoughts. *Int. J. Psychophysiol.* 29:1-11.

Llinas, R. R. 2001. *I of the Vortex: from Neurons to Self*. Cambridge MA: MIT Press.

Logan, G. D. 1992. Attention and preattention in theories of automaticity. *American Journal of Psychology* 105:317-339.

Mangan, B. 2001. Sensation's Ghost: The Non-Sensory "Fringe" of Consciousness. *Psyche* 7:http://psyche.cs.monash.edu.au/v7/psyche-7-18-mangan.html.

Marsden, C. D., J. C. Rothwell, and B. L. Day. 1984. The use of peripheral feedback in the control of movement. *Trends Neurosci.* 7:253-257.

Massimini, M., F. Ferrarelli, R. Huber, S. K. Esser, H. Singh, and G. Tononi. 2005. Breakdown of Cortical Effective Connectivity During Sleep. *Science* 309:2228-2232.

Mednick, S.A. 1962. The associative basis of creative process, *Psychol. Rev*. 69: 220–232.

Merker, Bjorn. 2005. The liabilities of mobility: A selection pressure for the transition to consciousness in animal evolution. *Consciousness and Cognition* 14: 89–114.

Minsky, M. 1985. *The Society of Mind*. New York: Simon and Schuster.

Miyake, A., and P. Shah. 1999. *Models of Working Memory*. Cambridge: Cambridge University Press.

Nadel, Lynn. 1992. Multiple memory systems: What and why. *J. Cogn. Neurosci.* 4: 179–188.

Nadel, L., and M. Moscovitch. 1997. Memory consolidation, retrograde amnesia and the hippocampal complex. *Curr. Opin. Neurobiol.* 7:217-227.

Ornstein, R. 1986. *Multimind*. Boston: Houghton Mifflin.Panksepp, J. 1998. *Affective Neuroscience*. Oxford: Oxford University Press.

Ramamurthy, U., S. D'Mello, and S. Franklin. 2004. Modified Sparse Distributed Memory as Transient Episodic Memory for Cognitive Software Agents. In *Proceedings of the International Conference on Sstems, Man and Cybernetics*. Piscataway, NJ: IEEE.

Ramamurthy, U., S. K. D'Mello, and S. Franklin. 2005. Role of Consciousness in Episodic Memory Processes. Ninth Conference of the Association for the Scientific Study of Consciousness - ASSC9. Pasadena; June 24 - June 27, 2005.

Ramamurthy, Uma, Sidney K D'Mello, and Stan Franklin. 2006. LIDA: A computational model of global workspace theory and developmental learning. In *BICS 2006: Brain Inspired Cognitive Systems*.

Reber, A. S., F. F. Walkenfeld, and R. Hernstadt. 1991. Implicit and explicit learning: individual differences and IQ. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 17:888-896.

Rubin, DC. 2006. The basic-systems model of episodic memory. *Perspectives in Psychological Science* 1: 277–311.

Rugg, M. D., and A. P. Yonelinas. 2003. Human recognition memory: a cognitive neuroscience perspective. *Trends in Cognitive Science* 7:313-319.

Schultheiss, O.C., Jones, N.M., Davis, A. Q., Kley, C. 2008. The role of implicit motivation in hot and cold goal pursuit: Effects on goal progress, goal rumination, and emotional well-being, *Journal of Research in Personality*, 42(4): 971-987.

Schooler, J. W. 2002). Representing consciousness: Dissociations between consciousness and meta-consciousness. Trends in Cognitive Science, 6, 339–344.

Seashore, C. 1967. *The Psychology of Music*. New York: Dover.

Sigman, M., and S. Dehaene. 2006. Dynamics of the Central Bottleneck: Dual-Task and Task Uncertainty. *PLoS Biol.* 4.

Sigman, M, A Jobert, D Lebihan, and S Dehaene. 2007. Parsing a sequence of brain activations at psychological times using fmri. *Neuroimage* 35, no. 2: 655-68.

Singer, Jerome L. (1978). Experimental studies of daydreaming and the stream of thought. In K. S. Pope & Jerome L. Singer (eds.), The Stream of Consciousness: Scientific Investigation Into the Flow of Experience. NY: Plenum Press.

Sloman, A. 1999. What Sort of Architecture is Required for a Human-like Agent? In *Foundations of Rational Agency*, ed. M. Wooldridge, and A. Rao. Dordrecht, Netherlands: Kluwer Academic Publishers.

Smallwood, J., Davies, J. B., Heim,D., Finnigan,F., Sudberry,M., O'Connor,R., and Obonsawin, M. (2004) Subjective experience and the attentional lapse: Task engagement and disengagement during sustained attention Consciousness and Cognition 13 (2004) 657–690.

Standing, L. 1973. Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology* 25:207-222.

Steinman, R. M., E. Kowler, and H. Collewijn. 1990. New directions for oculomotor research. *Vision Research* 30:1845-1864.

Stickgold, R., and M. P. Walker. 2005. Memory consolidation and reconsolidation: what is the role of sleep? *Trends Neurosci.* 28:408-415.

Taylor, J. G. 1999. *The Race for Consciousness*. Cambridge, MA: MIT Press.

Thompson, K. G., D. P. Hanes, N. P. Bichot, and J. D. Schall. 1996. Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *Journal of Neurophysiology* 76:4040-4055.

Tulving, E. 1983. *Elements of episodic memory*. Oxford: Clarendon Press.

Tulving, E. 1985. Memory and consciousness. *Canadian Psychology* 26:1-12.

Tversky, A., & Kahneman, D. (1973). Availability: a heuristic for judging frequency and probability. Cognitive Psychology 5, 207-232.

Uchida, N., A. Kepecs, and Z. F. Mainen. 2006. Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nature Reviews Neuroscience* 7:485-491.

VanRullen, R., and C. Koch. 2003. Is perception discrete or continuous? *Trends in Cognitive Science* 7:207-213.

Varela, F., J.-P. Lachaux, E. Rodriguez, and J. Martinerie. 2001. The brainweb: Phase synchronization and large-scale integration. *Nature Reviews Neuroscience* 2:229-239.

Willis, J., and A. Todorov. 2006. First Impressions: Making Up Your Mind After a 100-Ms Exposure to a Face. *Psychological Science* 17:592-599.