Modelling Memory and Learning Consistently from Psychology to Physiology

L. Andrew Coward
College of Engineering and Computer Science
Australian National University
Canberra, ACT 0200, Australia
andrew.coward@anu.edu.au

Abstract
Natural selection pressures have resulted in the physical resources of the brain being organized into modules that perform different general types of information processes. Each module is made up of submodules performing different information processes of the general type, and each submodule is made up of yet more detailed modules. At the highest level, modules correspond with major anatomical structures like the cortex, hippocampus, basal ganglia, cerebellum etc. In the cortex, for example, the more detailed modules include areas, columns, neurons, and a series of neuron substructures down to molecules. Any one memory or learning phenomenon requires many information processes performed by many different anatomical structures. However, the modular structure makes it possible to describe a memory phenomenon at a high (psychological) level in terms of the information processes performed by the major anatomical structures. The same phenomenon can be described at each level in a hierarchy of more detailed descriptions, in terms of the information processes performed by anatomical substructures. At higher levels, descriptions are approximate but can be mapped to more detailed, more precise descriptions as required down to neuron levels and below. The total information content of a high level description is small enough that it can be fully understood. Small parts of a high level phenomenon, when described at a more detailed level, also have a small enough information content to be understood. The information processes and resultant hierarchy of descriptions therefore make it possible to understand cognitive phenomena like episodic, semantic or working memory in terms of neuron processes via a series of intermediate levels of description.

1 Introduction

Understanding how the human brain supports higher cognitive phenomena like memory and learning cannot depend on models at the psychological level and other models at the physiological level with no clear connection between them. A hierarchy of descriptions of the same phenomenon from psychological to physiological is needed, with clear mapping between the levels (Coward and Sun 2004; Sun et al. 2005). At any one level, it must be possible to describe how at each point in time the observed situation causes the situation at the next point in time, in other words, descriptions must be causal.

The mapping between levels must be well understood. As in the physical sciences, the higher levels will be more approximate, but there must be clear understanding of when a more detailed level is necessary to achieve a given degree of quantitative accuracy (Coward and Sun 2007).

Such a hierarchy of causal descriptions requires consistent information models for components at each level of description. The information models make it possible to map precisely from a causal description on one level of detail to descriptions on other levels of detail.

For memory and learning, at the highest level a causal description in psychological terms is required. At a more detailed level, a causal description in terms of major anatomical structures, such as the cortex, thalamus, basal ganglia, cerebellum, etc., is needed that precisely maps into the psychological description. At an even more detailed level, a description in terms of components of the major anatomical structures (such as

cortical columns and subcortical nuclei) must map into the higher level descriptions. At an even more detailed level, a causal description in terms of neuron algorithms must map into higher anatomical descriptions.

Experience with the design of very complex electronic real-time systems indicates that a number of practical considerations place severe constraints on the architectures of such systems (Coward 2001). These practical considerations include resource limitations and the need to make changes and additions to some features without undesirable side effects on other features. "System architecture" means the way in which the information handling resources required to support system features are separated into subsystems. Each subsystem is individually customized for efficient performance of a different type of information recording, processing and/or communication, and there are limits on the type and degree of information exchange between subsystems. The separation between memory and processing and the sequential execution of information processes, often known as the von Neumann architecture, are important aspects of these architectural constraints.

Although there are minimal direct resemblances between such electronic systems and brains, natural selection results in a number of practical considerations that influence brain architectures in analogous ways (Coward 2001, 2005). If the brains of two species can learn the same set of behaviours, but the brain architecture of one of these species requires fewer resources, then the species with the more efficient architecture will have a significant natural selection advantage. If one species can learn new behaviours with less interference to previously learned behaviours than another species, then again the more effective species will have natural selection advantages.

It can be demonstrated theoretically (Coward 2001) that these and other practical considerations tend to constrain the architecture of any sufficiently complex learning system into some specific architectural forms, analogous with but qualitatively different from the von Neumann architecture. There is considerable evidence (Coward 2005) that the mammal brain has been constrained into these forms, known as the recommendation architecture.

As a result of these architectural forms, different general information models can be assigned to different major anatomical structures including the cortex, hippocampal system, thalamus, basal ganglia, amygdala, hypothalamus and cerebellum. More specific information models that support the general models can be assigned to substructures of these structures, such as areas and columns in the cortex; CA fields, dentate gyrus and associated cortices in the hippocampal system; and striatum, substantia nigra, globus pallidus, nucleus accumbens, etc. in the basal ganglia. Yet more specific supporting information models can be assigned to neurons, synapses and ion channels.

This hierarchy of consistent information models is the foundation for modelling memory phenomena consistently on a psychological level, on the level of major anatomical structures, on the level of more detailed structures and on the level of neuron physiology. Cognitive phenomena such as retrieval of episodic memories can be mapped into sequences of activities in major anatomical structures that result in the memory retrieval, an activity within one such structure (such as the cortex) can be mapped into sequences of activities in substructures of that structure (such as cortical columns), and an activity in a substructure can be mapped into a sequence of activities at a neuron level. If necessary, a neuron activity can be mapped into a sequence of chemical activities at the

synapse level, etc. The end result is understanding of the psychological phenomenon in terms of neurons.

2 The Recommendation Architecture Model

The information models for the major subsystems in the recommendation architecture are illustrated in figure 1, along with the anatomical structures of the mammal brain that correspond with each subsystem (Coward 1990, 2001, 2005, 2009a; Coward and Gedeon 2009). The primary separation is between a subsystem called clustering (corresponding with the cortex) that organizes the system resources used to define and detect conditions within the information available to the system and a subsystem called competition (corresponding with the basal ganglia and thalamus) that interprets each current condition detection as a set of recommendations in favour of many different behaviours, each with an individual weight, and implements the behaviour with the largest total weight across all current condition detections. Reward feedback results in adjustments to the weights that recommended recently implemented behaviours. However, such reward feedback cannot change condition definitions because such changes would interfere with all the other behaviours recommended by the same condition.

Clustering is organized as a modular hierarchy, with more detailed modules detecting sets of very similar conditions and higher level modules made up of groups of more detailed modules detecting larger sets of somewhat less similar conditions. The primary driving force that generates this hierarchy is the need to share the resources used to detect similar conditions. Individual modules detect conditions relevant to many different behaviours, and module outputs therefore have very complex behavioural meanings.

Competition is organized as a component hierarchy. The use of reward feedback within competition means that component outputs can only have very simple behavioural meanings, and components must correspond with one individual behaviour or with one general type of behaviour.

The information available to the cortex includes sensory inputs and inputs indicating internal activity of the brain (including the cortex). A condition is defined by a set of inputs and a specified state for each input. A condition is detected if a high proportion of the inputs that define it are in their specified state.

Two conditions are similar if there is significant overlap in the information defining them and/or they often tend to be present at the same time (i.e. in the same system input states). This definition of similarity implies that two similar conditions will tend to have similar behavioural implications. Rather than connecting every individual condition detection to the component hierarchy, resource economy can be achieved by organizing conditions into groups on the basis of similarity. The conditions making up the group are recorded in a module, and the module generates an output to the component hierarchy only if a significant subset of the conditions in the group is present. If necessary, sufficiently different subsets can be indicated by different outputs. The group of similar conditions defines the receptive field of the module.
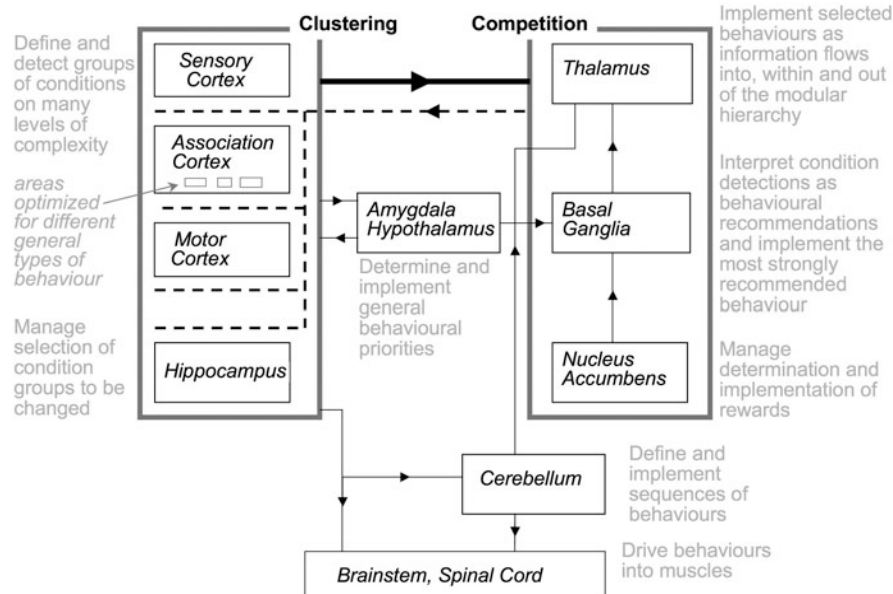
Figure 1 The recommendation architecture mapped into mammal brain anatomy. The primary separation is between a modular hierarchy (called clustering and corresponding with the cortex) that defines and detects conditions in the information available to the brain, and a component hierarchy (called competition and corresponding with the thalamus and basal ganglia) that at each point in time receives some of the conditions detected by clustering, interprets each condition as a set of recommendations in favour of a range of different behaviours, each with an individual weight, and implements the most strongly recommended behaviour. Most conditions must be defined heuristically, and within clustering there is a subsystem (corresponding with the hippocampal system) that determines when and where new conditions will be recorded. Positive or negative reward feedback following a behaviour can adjust the recently active weights that recommended the behaviour, but cannot change condition definitions without severe interference with all the other behaviours recommended by the same conditions. Detections of conditions indicating that a reward is appropriate are received by the nucleus accumbens, which applies changes to weights in the basal ganglia. Most behaviours are implemented by release of information flows into, out of or within clustering, and a separate subsystem (corresponding with the thalamus) is required to efficiently manage such information releases. Conditions indicating the appropriateness of general types of behaviours (e.g. aggressive, fearful, food seeking, etc.) are provided to structures (corresponding with the amygdala and hypothalamus) that modulate the relative probability of such behaviour types. Frequently required sequences of behaviours that need to be executed rapidly and accurately are recorded and executed by the cerebellum.

Conditions (and receptive fields) can be defined on many different levels of complexity, where the complexity of a condition is the total number of raw system inputs (including all duplicates) that contribute to the condition, either directly or via intermediate conditions. Receptive fields on different levels of complexity will tend to be effective in discriminating between different types of circumstances with different types of behavioural implications. For example, relatively simple receptive fields will be able to discriminate between different visual features (e.g. tail, wing and tooth). More complex receptive fields will be able to discriminate between different categories of visual object (e.g. cat, dog and bird). Yet more complex receptive fields will be able to discriminate between different types of groups of visual objects (e.g. cat chased by dog, cat confronting dog and cat avoiding dog). Simple receptive fields will be elements in the definitions of more complex receptive fields.

In a learning system, most conditions and receptive fields must be defined heuristically from system experience. Hence the information model for the cortex as a whole includes definition and detection of receptive fields at different levels of complexity within the information available to the brain.

In the basal ganglia, high-level components correspond with general types of behaviour, more detailed components within a high level component correspond with more specific behaviours of the general type. A receptive field detection by the cortex is communicated to a large number of components in the basal ganglia. Each component interprets the detection as a recommendation in favour of its corresponding behaviour, with a weight that is different for each component. The information model for the basal ganglia is therefore interpretation of cortical receptive field detections as behavioural recommendations and implementation of the most strongly recommended behaviours.

In general, a behaviour will be implemented by an information flow into, within or out of the cortex. There is therefore a requirement for a subsystem that provides coordinated management of all such information flows. This subsystem requires information on current cortical activity and information on currently selected behaviours. In the mammal brain, the thalamus corresponds with this subsystem. Thus the thalamus receives information from the cortex and gates information flows into the cortex (attention behaviours), out of the cortex (e.g. from the motor cortex to drive motor behaviours) and within the cortex (e.g. various cognitive behaviours).

If a behaviour is implemented, it will have consequences. If these consequences are good, the probability of the behaviour being selected in the future can be increased, if the consequences are bad, the probability should be decreased. A reward subsystem is therefore required that receives indications of positive and negative consequences of behaviours and changes recently active weights that led to the selection of recent behaviours. The nucleus accumbens within the basal ganglia corresponds with this subsystem. Because components in the component hierarchy correspond with behaviours, appropriate targetting of reward feedback is straightforward.

Indications of the positive or negative consequences of behaviours will often be receptive field detections by the cortex. For example, the anterior cingulate cortex detects receptive fields that correlate with the presence of error conditions (Kiehl et al. 2000). The anterior cingulate cortex projects to the nucleus accumbens (Devinsky et al. 1995), where such receptive field detections are interpreted as recommendations in favour of negative rewards. If the recommendations are accepted (i.e. if there is adequate total weight into the nucleus accumbens), then the recently accepted behavioural weights elsewhere in the basal ganglia will be weakened.

Any one receptive field detection recommends a wide range of different behaviours, and reward feedback follows an individual behaviour. Hence changes to receptive field definitions as a result of reward feedback would damage the integrity of the recommendation weights in favour of all the other behaviours. In general, any major change to a receptive field will damage the integrity of its associated behavioural meanings. As a result, receptive field changes must be strongly constrained. To a first approximation, a receptive field can be expanded slightly by addition of similar conditions, but previously added conditions cannot be changed or deleted. One exception to this rule is that if a condition is added and does not occur again within a significant period of time, it can probably be removed again. Another exception is that if a receptive

field ceases to occur for some long period of time (perhaps because the source of its inputs has been damaged), it can probably be removed and its resources reassigned to another receptive field.

An implication of this constraint on receptive field changes is that individual fields cannot be guided to correspond with unambiguous cognitive circumstances such as object categories. Any one receptive field may be detected within some instances of many different such object categories and will therefore have recommendation strengths in favour of behaviours appropriate to all of the categories (such as naming them). However, an array of such fields can be evolved so that it discriminates between such categories. In other words, the predominant recommendation strength across all the receptive fields detected within an instance of any category will be in favour of behaviours appropriate for that category, even though no one receptive field corresponds with any one category (Coward 2001, 2005).

There is another important implication of the stability of receptive fields and their associated behavioural meanings. In response to an input state, a number of receptive fields will be detected, contributing their behavioural recommendation strengths. However, there are some potentially relevant recommendation strengths that are associated with receptive fields not actually being detected. For example, suppose there is a receptive field that is not currently being detected, but that has often been detected in the past when many of the currently detected receptive fields were also detected. Given this past consistency, there could be behavioural value in generating receptive field detections on the basis of frequent past simultaneous activity with currently detected receptive fields. Such value could be accessed by each receptive field also having recommendation strengths in favour of indirect activation of any receptive fields that have often been active in the past at the same time. Behavioural value could also exist for receptive field activations on the basis of recent simultaneous activity and also on the basis of simultaneous receptive field expansion. Furthermore, there could also be behavioural value in receptive field activations on the basis of past activity just after the activity of currently detected receptive fields, where the temporally correlated activity could be frequent, recent or receptive field expansion based.

If unrestricted, such indirect activations would generate a huge amount of activity, and they must therefore be behaviours that compete with other behaviours for acceptance. These indirect activation mechanisms are the information basis for a wide range of memory phenomena.

In order to achieve a high integrity behaviour, there must be an adequate range of recommendations available. To achieve an adequate range, there must be at least a minimum number of receptive field detections in response to every overall system input state. If this minimum is not reached, some receptive fields must expand in order to extend the range of available recommendations. However, to minimize changes to receptive fields, those requiring the least degree of expansion must be identified.

There is therefore a requirement for a further major subsystem that determines whether receptive field expansions are required, and if so identifies the modules in which the least such expansion will be required and drives expansion in those modules. In the mammal brain, the hippocampal system corresponds with this resource management subsystem. The information model for the hippocampal system is therefore determination

of when and where in the cortex receptive field expansions are appropriate and driving the required expansions.

There is an additional requirement to record frequently used sequences of behaviours to ensure their rapid and accurate execution whenever required. In the mammal brain, the cerebellum corresponds with this subsystem.

There is a requirement for a subsystem that modulates the relative probabilities of different general types of behaviours. Such general types could include aggressive, fearful, food-seeking, etc. There are two ways in which relative probabilities could be affected. One is to change the relative arousal of components in the component hierarchy corresponding with the behaviour types. The other is to change the probability of recommendations of the general type being generated. To economize on resources, most receptive field detections must be able to recommend any type of behaviour. However, if there are substantial behavioural advantages in separate receptive field definitions at some levels of complexity for different behavioural types, the advantages might outweigh the resource costs. Hence there could be some modules within the cortex that tend to recommend one general type of behaviour, and temporarily broadening their receptive fields would increase the probability of such behaviours being selected.

Finally, there are two dynamic considerations. Efficient use of receptive field detection resources requires that the same set of resources must be able to detect receptive fields at one level of complexity simultaneously within multiple different sensory circumstances, keeping the different detections separate until it is appropriate to combine them in a controlled fashion. For example, if two dogs are chasing a squirrel, receptive fields must be detected within each dog and also the squirrel. There must be a set of receptive fields at a level of complexity that is effective for discriminating between objects. Resource economy dictates that each object will result in receptive field detections which are subsets of the same set, generally with some overlap between the subsets. Once receptive fields have been detected within all the individual objects, the detections must be combined to detect more complex receptive fields that can discriminate between different types of groups of objects. The problem is that the sets of simpler receptive fields must be kept completely separate (e.g. one activated set of receptive fields must not correspond with the head of a dog on the body of a squirrel) but must all be active at the same time for them to be combined to detect the more complex receptive fields (two dogs chasing a squirrel). As will be discussed in more detail later, one role of the gamma band frequency in the EEG is to maintain a separation between receptive fields detected within different objects but using the same neural resources.

The second dynamic consideration is that because there is behavioural value in indirect activation of receptive fields on the basis of activity shortly after currently active fields, there is an analogous requirement for simultaneous activity in the same resources corresponding with a sequence of points in time. This simultaneous activity is required so that the appropriate links supporting future indirect activations can be established. As discussed later, one role of the theta band frequency in the EEG is to maintain separation between receptive field detections corresponding with different points in time.

3. Review of Experimental Data Literature
Tulving (1984, 1985) and later Schacter and Tulving (1994) developed an approach to classifying memory and learning phenomena. They suggested there were three criteria

that could be used to identify a memory system in the brain, which was separate from other memory systems. These criteria are the existence of a group of memory tasks which have some common characteristics, the existence of a list of features different from the list for any other system and the existence of multiple dissociations between any two systems. A dissociation is a way in which a similar manipulation of tasks performed by different systems produces different effects on the performance. An important category of dissociations is observations of patients with brain damage that affects one type of memory and not another.

On the basis of a wide range of evidence, Schacter and Tulving (1994) proposed that there are five independent memory systems: semantic memory; episodic memory; priming memory; procedural memory and working memory. Semantic memory is the memory for facts and the meaning of words, without recall of the context in which those facts or words were learned. Episodic memory is the memory for events, including autobiographical memory for events with personal involvement. Semantic and episodic memories are together known as declarative memory because they are consciously accessible and can be described verbally. Priming memory is the ability to make use of experiences in the recent past to enhance behaviour, even when there is no conscious awareness or memory of those experiences. Procedural memory is the ability to learn skills, including motor skills. Working memory is the ability to maintain direct access to multiple objects, so that information derived from the objects is immediately available to influence behaviour.

3.1 Semantic Memory

Semantic memory is defined as the ability to recall a wide range of organized information including facts and word meanings (Tulving 1972). The typical experimental test of semantic memory is category verification, where a subject is presented with paired category names and category (e.g. mammal, monkey or mammal and pigeon) and asked to identify if the pairing is correct. Identification speed is slightly slower for non-typical instances (e.g. mammal-bat) than for typical or incorrect pairings (Rips et al. 1973).

Functional neuroimaging indicates that semantic knowledge is encoded in many different cortical areas, especially the posterior temporal and frontal cortices, with the areas activated during semantic memory tasks generally being those also active during sensory or motor processing (Martin 2007). However, damage to the most anterior portions of the temporal cortices results in general loss of semantic memory capabilities, although this area does not show strong activation during semantic memory tasks (Rogers et al. 2006). There appears to be no consistent evidence for cortical area specialization for semantic domain (e.g. natural or man-made objects) or category (e.g. animals, fruit, tools and vehicles) although there is animal specific activity in the left anterio-medial temporal pole and tool specific activity in the left posterior middle temporal gyrus which appears in a subset of experiments with lower statistical confidence (Devlin et al. 2002).

3.2 Episodic Memory

In the laboratory, episodic memory is measured by both recognition and recall experiments. In recognition experiments, subjects are shown a set of novel objects and later shown a mixture of further novel objects and objects from the earlier set and asked to identify objects seen before.With photographs, subjects have a remarkable high

capability to identify previously perceived objects (Standing et al. 1970). In recall experiments, subjects are asked to describe one past event. To trigger the recall, subjects are given a word (Robinson 1976) or group of words (Crovitz and Schiffman 1974). Such experiments are often used to measure the past time period for which episodic memory retrieval has been degraded and the degree of degradation (Kensinger et al. 2001). 327

Observation of the severity of retrograde amnesia in amnesic patients with damage to their hippocampal system indicates graduations between semantic and episodic memory (Nadel and Moscovitch 1997). In such patients, the most severe amnesia is for personal autobiographic memories. Amnesia for personal information, public events and persons is less severe, and amnesia for words and general facts is often minimal.

Functional neuroimaging of the brain during episodic memory recall indicates that there is strong activity in the prefrontal cortex during episodic recall (Fletcher et al. 1997) and also in visual and visual association areas (Addis et al. 2007). There is somewhat weaker activity in the hippocampal system (Fletcher et al. 1997). Strong cerebellar activity is observed during episodic memory retrieval (Fliessbach et al. 2007).

3.3 Procedural Memory

Experimental tests of procedural memory are generally confined to simple learning tasks and are focussed on clarifying the distinction between procedural and other types of memory, in particular declarative memory (meaning both semantic and episodic). Typical investigations investigate the ability of patients who have lost the ability to create new semantic and episodic memories to learn simple procedural skills. For example, a wide range of such amnesic patients showed the ability to learn to read words reflected in a mirror at the same rate as normal controls and retain the skill for at least 3 months, despite failing to recall any familiarity with the task at the start of each session (Cohen and Squire 1981).

For more complex skills, declarative knowledge speeds up the learning process (Mathews et al. 1989; Sun et al. 1996), and declarative memory is required for high levels of procedural skill performance (Mathews et al. 1989). However, there can be inconsistencies between procedural and declarative knowledge.When highly skilled subjects describe their skill, the descriptions often correspond with beginner methods rather than actual methods, and generating a verbal description can result in reversion to the less effective beginner method (Bainbridge 1977; Berry 1987).

There is a range of evidence derived from the cognitive deficits associated with degeneration of the basal ganglia indicating that this structure plays an important role in procedural learning. The symptoms of Parkinson's disease include difficulty with voluntary movement and with initiation of movement and in general slowness of movement. The observed physical deficit (Jankovic 2008) is degeneration of dopaminergic neurons in the substantia nigra compacta (SNc) nucleus of the basal ganglia. The major symptom of Huntingdon's disease is the intrusion of irregular, unpredictable, purposeless, rapid movements that flow randomly from one body part to another (Berardelli et al. 1999). The observed physical deficit is loss of striatal cells that project into the indirect pathway (Starr et al. 2008).

In addition, the cerebellum plays an important role in procedural learning (Torriero et al. 2007), although it also plays a role in a wide range of higher cognitive processes

(Leiner et al. 1993) including semantic (Devlin et al. 2002), episodic and working memory (Cabeza et al. 2002).

3.4 Working Memory

Working memory refers to the number of different objects that can be maintained active in the brain at the same time. A typical working memory experiment is list recall. Subjects are shown a sequence of objects and immediately afterwards asked to list all the objects in any order. Normal subjects can fully recall sequences of seven (plus or minus two) random digits, but only four or five random words or letters (McCarthy and Warrington 1990). Miller (1956) suggested that the limit of seven plus or minus two is a fundamental information processing limit. Cowan (2000) argued that seven is an overestimate because subjects were able to rehearse and/or chunk items and proposed that the true information limit is close to four items, based on a wide range of observations. One key observation is that there are performance discontinuities around the number four, with errorless performance below four and sharp increases in errors above four (Mandler and Shebo 1982).

In a variation of the list recall test, the number of items that subjects can report on longer lists is measured. Typically, there is enhanced recall for the first few items on the list and enhanced recall for the last few items (Baddeley 2000). A brief delay occupied with another task eliminates the recency effect but has much less influence on the primacy effect (Glanzer 1972). List recall capability is greater if there is a semantic connection between objects on the list. Recall of meaningful sentences can extend to 16 words or more, but recall for random words is limited to four or five (Baddeley et al. 1987).

Neuroimaging indicates considerable overlap in the cortical areas active during working memory and declarative memory tasks. For example, Cabeza et al. (2002) used fMRI to compare brain activity during an episodic memory task (recalling if a word was on a list of 40 words studied much earlier) with a working memory task (recalling if a word was in a list of four words presented 15 s earlier). They found that the cerebellum and left dorsolateral cortex areas were active during both tasks, bilateral anterior and ventrolateral cortex areas were more active during episodic retrieval and Broca's area and bilateral posterior/dorsal areas were more active during working memory retrieval. A patient with damage to the left parietal lobe showed a deficit in working memory but unaffected declarative memory (Warrington and Shallice 1969; Shallice and Warrington 1970).

3.5 Priming Memory

Priming memory is a short-term effect of exposure to a stimulus on the response to a similar later stimulus. Priming memory decays rapidly over a period of minutes, then more slowly over periods of hours and days. One experimental test of priming is word stem completion, in which subjects are asked to complete each of a list of three letter word stems with the first English word that comes to mind. A stem is the first three letters of a word, and in such experiments there are typically about ten possible completions. Previous study of the word increases the probability of the word being generated, provided the study is less than about 2 h prior to test (Graf et al. 1984). Amnesic patients

with no ability to create new declarative memories have priming memory that is the same as and decays at the same rate as in normal subjects (Graf et al. 1984).

Another important priming memory experiment is tachistoscopic image recognition, when a subject is shown a sequence of brief (<100ms) presentations of images of objects, each image separated by a masking pattern to prevent retinal, etc. afterimages. In this situation, few images can be accurately identified. However, if later there is a repeat exposure to the same image, identification accuracy increases considerably (Bar and Biederman 1998; Badgaiyan 2000).

It has been argued that unconscious priming involves different mechanisms from when the earlier stimulus is consciously used. However, McBride et al. (2001) presented evidence that similar mechanisms operate in both cases. Their experiments used word fragments, which were words from which two to four letters had been replaced by spaces. Subjects were given word lists to study. After a measured period of time, they were given word fragments and asked to complete the word. Two types of experiment were performed. In one, the subjects were asked to complete the fragment with a word studied earlier (i.e. a conscious approach). In the other, they were asked to complete the fragment with the first word that came to mind (i.e. an unconscious approach). The proportion of completions with studied words declined with time, rapidly in the first 10 min, but both performance and rate of decline were the same with both conscious and unconscious instructions.

3.6 Dissociations Indicating Separate Memory Systems
The evidence for the different memory systems is extensively discussed in Schacter and Tulving (1994). Any theory of memory and learning at anatomical and physiological levels must provide an account for this evidence.

Some of the most striking evidence for the separation of semantic and episodic memory systems from each other and from other memory types comes from observations of patients with damage to their hippocampal systems. In the 1950s, a number of patients had experimental surgery to treat intractable epilepsy. The surgery involved sectioning significant parts of their hippocampal systems, and although successful in reducing the frequency of epileptic seizures, it had some drastic side effects on their memory capabilities (Scoville and Milner 1957). One of these patients, HM, was extensively studied until his death in 2008.

HM lost all capability to acquire new semantic or episodic memories. He could learn no new facts or words and could recall no events after his surgery (Scoville and Milner 1957; Corkin 2002). However, he retained normal working memory (Wickelgren 1968) demonstrating a dissociation between working memory and declarative memory. He retained skills learned prior to surgery, including speech and reasoning skills (Scoville andMilner 1957). He could still acquire simple motor skills such as learning mirror writing, showing steady improvement over a number of sessions, even though at each session he had no memory of attempting the task before (Gabrieli et al. 1993), demonstrating a dissociation between declarative and procedural memory. His priming memory was retained (Gabrieli et al. 1990). In addition, his ability to access episodic memories for 11 years prior to his surgery was also impaired (Sagar et al. 1985), but his semantic memory for word meanings learned in the same 11-year period was retained

(Kensinger et al. 2001), demonstrating a dissociation between episodic and semantic memory.

Evidence for the separation between priming and procedural memories comes from the study of Huntington's Syndrome patients. Such patients are characterized by damage to the basal ganglia and exhibit severe deficits in motor skill learning, but their priming memory appears intact (Heindel et al. 1989).

As mentioned earlier, patients have been observed to exhibit deficits in working memory, with no apparent deficit in declarative memory (Warrington et al. 1971). In these patients, performance in an immediate memory span test in which they were presented with strings of one to four digits, letters or words revealed good recall for one item strings but well below normal recall for strings with more than one item. However, performance in recall of a short story showed performance slightly better than for normal subjects.

4 Other Modelling Approaches

There is a very large number of attempts to model memory phenomena, ranging from attempts to prove that synaptic weight changes are always present during learning (Martin and Morris 2002) to high-level psychological models such as Baddeley's working memory model that uses subsystems like phonological memory and central executive with no attempt to map into physiology. In general, previous models focus on one or two levels of description and do not present a consistent hierarchy of descriptions with information models on each level which can be mapped between all levels from psychology to physiology.

The five system memory model proposed by Schacter and Tulving (1994) and described earlier is in fact a psychological level model with some implications for high level anatomy. As discussed earlier, there have been successful attempts to map the model into major anatomical structures by the evidence from deficits resulting from local brain damage (Schacter and Tulving 1994) and various imaging techniques including fMRI and PET. For example, Devlin et al. (2002) and Rogers et al. (2006) analyze the brain regions active during semantic memory processes. Kassubek et al. (2001) investigated brain regions active during procedural memory processes. Fletcher et al. (1997) and Addis et al. (2007) have investigated the brain regions active during episodic memory processes. These investigations demonstrate differences in the cortical areas active during different types of memory and are valuable high level descriptions, but do not provide information models for the processes that can be mapped into deeper level descriptions.

Another type of approach has been the development of phenomenological models like Baddeley's working memory model (1986) and the spreading activation model for semantic memory (Collins and Loftus 1975). These models attempt information models for the phenomena at high level but do not offer any mapping to more detailed levels. The ACT-R model (Anderson 1996) offers a detailed information model which can, for example, model working memory in a fair amount of quantitative detail (Lovett et al. 1999). However, although there have been attempts to map ACT-R to fMRI imaging, the ACT-R information models do not provide any mapping into information models for more detailed anatomical or physiological structures.

Another extensive modelling approach to memory is that of Hasselmo and his collaborators. For example, they have proposed a model for the operation of working memory and episodic memory that uses reinforcement learning and attempts to account for a range of observations on rats (Zilli and Hasselmo 2007). This model postulates the existence of a number of buffers which can be in states reflecting current sensory inputs or a range of past sensory inputs that have been recorded. A key aspect of the model is the concept of actions which can be taken on the memory systems themselves in addition to actions on the external environment. There are some analogies between these proposed self actions and the indirect activation of receptive field information in the recommendation architecture. An implementation of their model has been described (Zilli and Hasselmo 2007), but the implementation does not provide mapping into plausible physiology. Furthermore, the model does not contain the recommendation architecture separation between clustering (i.e. condition definition and detection) and competition (i.e. reinforcement learning-based interpretation of conditions into behavioural recommendations). As a result, the model would have problems scaling up to learn complex combinations of behaviours.

At the neurophysiological level there have been numerous proposals that synaptic plasticity supports memory (e.g. Martin and Morris 2002). However, although these proposals are based on experimental evidence that long-term changes to synaptic weights are associated with learning, they do not offer neuron level information models that can be mapped (through intermediate anatomical levels) into, for example, semantic and episodic memory.

Another important set of models are those which attempt to understand the role of different EEG frequencies in memory. The beta frequency (12–30Hz), gamma frequency (30–80Hz) and the theta frequency (5–12Hz) occur throughout the neocortex and hippocampus. These frequencies appear as modulations placed upon the firing of pyramidal neurons and are probably managed by interneuron activity (Whittington and Traub 2003). A number of proposals have been made that these frequencies play various roles in memory.

For example, Hasselmo et al. (2002) proposed that different functions are supported in different phases of the theta frequency in the hippocampus. In one part of a theta cycle, associations between sensory events are learned. In the other part of the cycle, previously learned associations are retrieved. As a result, the theta rhythm plays an important role in the reversal of previously learned associations (e.g. when the physical location of a reward changes). This model has been simulated with a considerable degree of physiological detail (Cutsuridis et al. 2008, 2010; Cutsuridis and Wenneckers 2009). However, the model focusses on the hippocampus and does not address the role of the cortex in these memory functions, other than the general view that it is the long term storage location for declarative memory, with the hippocampus providing intermediate term storage. How declarative memories are transferred from hippocampus to cortex is not addressed.

A number of workers have suggested that the gamma frequency binds together the activity of neurons representing the features of the stimulus that is the focus of attention (e.g. Engel and Singer 2001). To this proposal has been added the idea that the gamma frequency is important for the formation of declarative memories (Axmacher et al. 2006).

There is a relationship between these proposals and the recommendation architecture requirement for separation between different populations of receptive field detections correspondingwith different cognitive stimuli but within the same neural resources. This relationship will be discussed below. However, the various dynamic models do not provide causal descriptions of higher level cognitive phenomena which can be mapped into the dynamic models.

Previously proposed memory models can model memory phenomena at one or two levels of detail, but do not provide a hierarchy of information models which can map consistently from psychological phenomena to neuron physiology.

The primary focus of this chapter is on memory and learning phenomena. However, the recommendation architecture is a general cognitive architecture that can provide descriptions of a wide range of higher cognitive phenomena in terms of brain anatomy and physiology (Coward 2005). A comparison between the recommendation architecture approach and a wide range of alternative cognitive architectures including Haikonen's neural architecture, Baar's global workspace, virtual machine models, simulation models, kernel architectures and forward models has been performed (Coward and Gedeon 2009). This study concluded that the alternative approaches do not take adequate account of natural selection pressures on brain resources, and the mapping between higher cognition and detailed anatomy and physiology is more plausible for the recommendation architecture.

5 Brain Anatomy and the Recommendation Architecture Model

The RA cognitive architecture maps the information model for each major anatomical structure into more detailed models for its substructures and so on all the way down to neuron physiology. The architecture explains how the more detailed models are implemented physiologically, and how the detailed models interact to support higher level models, up to descriptions of memory on a psychological level.

5.1 Cortical Structure

The cortex is a 3–4mm thick sheet of tissue, with an area in adult humans of about 2,600 cm^2. The cortical sheet is organized into six layers, with the layers differing in cell type, size and density and in intralayer and interlayer connectivity. Six layers are generally prominent, but sublayers are often visible, and sometimes a major layer can be absent. The cortex is organized perpendicular to the sheet into groups of about 100 cells linked across the layers. These groups, called minicolumns, are produced by the cortical growth process, each minicolumn from a small set of progenitor cells. Cortical columns are much larger vertical structures, perhaps formed by binding together a number of minicolumns. Cortical columns vary from 300 to 600 μm in diameter and are distinguished by similarity of the receptive fields of all their principal neurons and by common short-range horizontal connections (Mountcastle 1997).

The cortical sheet is separated into at least 50 areas (Brodmann 1908; Petrides and Pandya 1999, 2002; Morasan et al. 2001). These areas differ in the cell types, sizes and densities observed in each layer and sublayer, in the concentration of various neurochemicals such as the neurofilament protein that influences neuron size and shape, in the relative degree of myelinization of each layer and sublayer, in the interconnectivity with adjacent layers and in the interconnectivity with other areas.

5.2 Cortical Information Models

As discussed earlier, the information model for the cortex is definition and detection of receptive fields on many different levels of complexity within the information available to the brain. The number of detections must reach at least a minimum but not an excessive level in response to every overall brain input state. A receptive field is defined by a set of similar information conditions, with the receptive field being detected if a significant subset of the conditions is detected.

5.2.1 Information Model for a Cortical Area

The information model for an area is definition of receptive fields within one range of complexity and detection of the most significant receptive fields in each input state. These receptive fields must adequately address five requirements. The first requirement is that, in order to preserve previously acquired behavioural meanings, receptive fields can expand to detect additional, similar conditions but with some tightly restricted exceptions cannot change or discard previously added conditions. The second requirement is that, in response to any brain input state that results in inputs to the area, the number of receptive field detections must reach at least a minimum level. This second requirement is to ensure that, as discussed earlier, enough alternative behavioural recommendations are available to generate a high integrity behavioural selection. If detections are below the minimum, some receptive fields will expand until the minimum is reached.
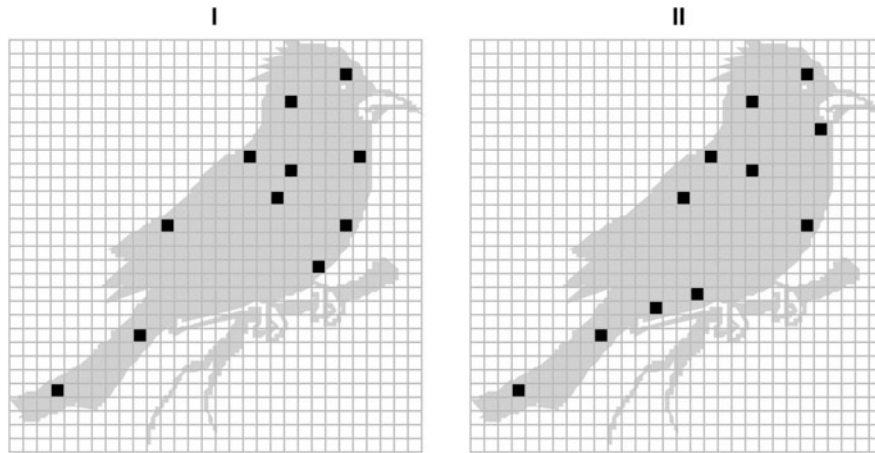

Figure 2 The concept of conditions and condition similarity. A visual domain of 30 x 30 pixels is defined. A visual object is present in the domain. A condition is defined by a set of elements (in this case pixels), with a state specified for each element of the set. Such conditions are illustrated in a and b. If a visual input results in a high proportion of the pixels in the set being in their specified state, the condition is present. Two conditions are similar if a high proportion of the elements and states that define them are the same and/or often occur in the same input states. Conditions I and II are very similar by this definition, since seven of their elements are the same, and all of the elements occur in the same input state. With this definition of similarity, similar conditions will tend to have similar behavioural implications (such as recommending saying "that is a bird")

The third requirement is that, in response to any brain input state that results in inputs to the area, the number of receptive field detections does not reach an excessive level. The fourth requirement is that, to conserve resources, two receptive fields should not be

detected consistently in the same input states. In other words, receptive fields must be as orthogonal as possible, or as statistically independent as possible. The fifth requirement is that the set of receptive fields programmed in one area must be able to discriminate between different perceptual circumstances, whenever such a difference implies that a different behavioural response is appropriate. The type of perceptual circumstance is different for different areas, the receptive fields in one area discriminate between visual features, in another between object categories, in yet another between different types of groups of objects, etc. as discussed earlier.
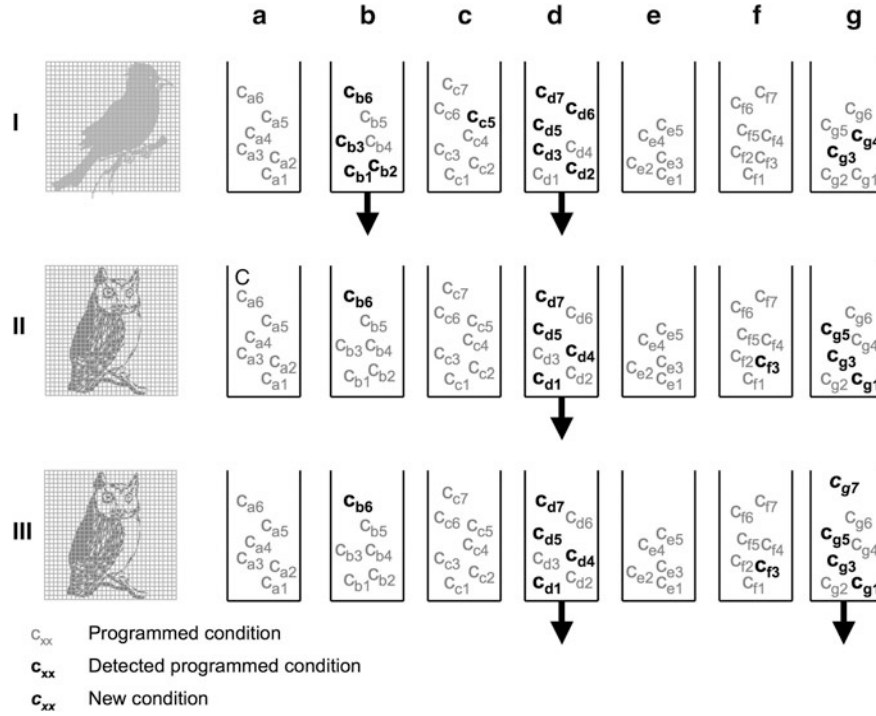
Figure 3 The concept of receptive field expansion management. In the illustration there are six modules, each module having a receptive field defined by a set of similar conditions as described in figure 2. If a significant proportion of the conditions programmed in a module are detected, the module receptive field is detected and the module produces an output. The array of modules is required to detect at least two receptive fields in every input state. In scenario I, the visual input state results in detection of a significant proportion of conditions in modules b and d, corresponding with receptive field detection and resulting in outputs from those modules. The level of condition detection in other modules is relatively low. In scenario II, the visual input state is less familiar, and initially only module d has enough condition detection for receptive field detection. Because this is less than the minimum, the module with the highest degree of condition detection below the proportion for receptive field detection (i.e. module g) is selected, and additional conditions actually present in the current input are added to the module. Such additions expand the receptive field of the module and result in enough condition detections for receptive field detection and module output (scenario III). Selection of the module without receptive field detection but with the highest degree of condition detection is equivalent to selecting the module requiring the least receptive field expansion

The conceptual process for definition of receptive fields can be understood by consideration of figures 2 and 3. In figures, there is a visual input domain which is an array of pixels. Visual inputs appear within the domain. A condition is defined by a set of elements (in this case pixels) and a state for each element, the condition is detected if a

16

high proportion of the elements are in their defined state. As illustrated in figure 2, two conditions are similar if their elements are the same and/or often occur in the same input states. The modules illustrated in figure 3 have receptive fields defined by groups of similar conditions, with the receptive field being detected if a high proportion of the conditions is detected. An array of modules such as the one illustrated in figure 3 detects different receptive fields, all in one range of complexity. Such an array is required to detect at least a minimum number of receptive fields in every input state. If the minimum is not achieved, modules that are detecting a significant number of conditions but less than the proportion for receptive field detection add conditions until the proportion for receptive field detection is reached. These conditions are selected (by a biased random process to be discussed later) from the large number of conditions that are present in the current input state. Receptive field expansions therefore occur in those modules for which the smallest expansion is required, causing the least damage to existing behavioural meanings of the receptive fields.

Arrays of modules are arranged in a sequence, with the conditions detected in one array being combinations of the conditions detected in the preceding array. Hence there is a steady increase in receptive field complexity down the sequence, where complexity can be defined as the number of rawsensory inputs (in the conceptual example pixels) that must be present in their appropriate state for the receptive field to be detected.
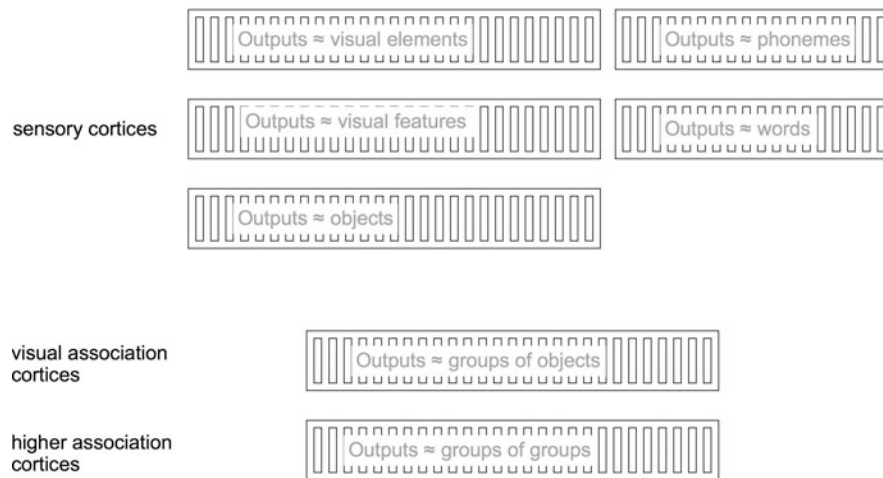

Figure 4 Sequence of arrays of modules in which the conditions detected in one array are combinations of receptive field detections by modules in the preceding array. Few receptive fields will correspond with cognitively unambiguous circumstances, but at the appropriate level of receptive field complexity, the sets of receptive fields detected within instances of one category will be sufficiently different from the sets detected within instances of any other category that the currently detected set can effectively guide selection of a behaviour appropriate to current category. In other words, the array of columns will be able to discriminate between different categories. Receptive fields on different levels of complexity will be able to discriminate effectively between different types of cognitive categories (features, objects, groups of objects, etc.). Multiple levels of receptive field complexity may be relevant to, for example, discrimination between different types of groups of objects

It may be possible to describe receptive fields in areas close to sensory input in sensory terms [such as the class of visual shapes that contain the receptive field, see Tanaka (2003)]. However, for higher areas, receptive fields can best be understood

simply as groups of receptive fields in other areas that have tended to be detected at similar times in the past. The receptive field in one area receives inputs from the group of receptive fields in other areas and is detected if a high proportion of the group is present.

The mechanism for definition of receptive fields means that in general one receptive field will not correspond with one unambiguous cognitive circumstance such as an object category. In other words, no single receptive field will always be detected in all instances of one object category and in no instances of any other category. However, an array can be evolved to discriminate between types of cognitive circumstance, with arrays with different receptive field complexities being effective for discriminating between different types as illustrated in figure 4.
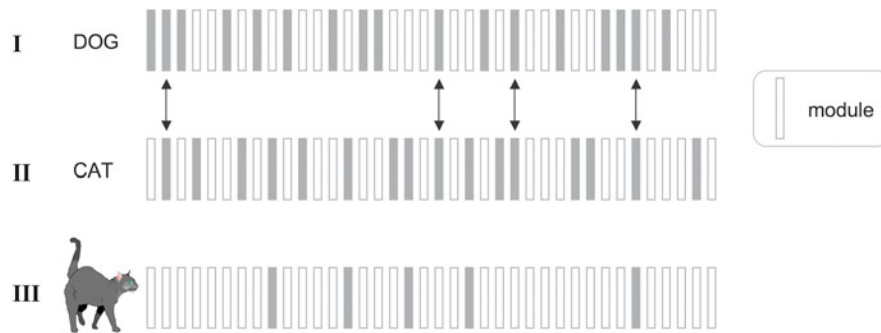

Figure 5 An array of modules detecting receptive fields at a level of complexity effective for discriminating between different visual object categories like DOG and CAT. In I, the modules that detect their receptive fields in some instances of DOG are shaded. A different subset of these shaded modules will detect its receptive field in each actual instance of DOG, and each shaded module will have some recommendation strengths (in subcortical structures) in favour of DOG-appropriate behaviours such as saying "that is a dog". In II, the modules in the same array that detect their receptive fields in some instances of CAT are shaded. There is some overlap between the CAT and DOG modules illustrated in I and II, reflecting the existence of some visual similarities between cats and dogs. Receptive fields detected in some instances of both CAT and DOG will have recommendation strengths in favour of both DOG- and CAT-appropriate behaviours. All the modules will also have recommendation strengths in favour of behaviours appropriate to many other categories of object. In III, the modules in the same array that detect their receptive fields in one actual CAT instance are shaded. Although some modules with DOG-appropriate recommendation strengths are active, the predominant recommendation strength across all receptive fields will be CAT-appropriate. The module array can thus discriminate between dogs and cats.

The meaning of discrimination can be understood from figure 5. The modules in the array illustrated in the figure detect receptive fields within a range of complexity appropriate for discriminating between visual objects. Each receptive field corresponds with some visual circumstance which may occur in some instances of many different object categories. Thus, for example, one module may detect its receptive field in some instances of the category DOG, but also in some instances of the category CAT. In general, one receptive field will be detected in some instances of many different categories. There are no modules that detect their receptive fields in all DOG instances and in no instances of other categories, and individual modules are therefore cognitively ambiguous. However, the set of modules that detect their receptive fields in response to an actual CAT instance will have a predominant recommendation strength in favour of CAT-appropriate behaviours.

Module receptive field definition is not directly guided by the existence of cognitive categories, and an array might lack discrimination in some cases. A discrimination problem and the mechanism for resolving it are illustrated in figure 7. In the figure, because of a visual similarity an instance of a dog and an instance of a cat activate the same receptive fields. In other words, the array of receptive fields cannot discriminate between the two instances. The effect will be that the same behaviours will be recommended for both objects. Suppose that the dog was seen first, a DOG-appropriate behaviour was performed and rewarded, and the set of columns now have a strong overall recommendation strength in favour of repeating that behaviour in the future. When the cat is seen, there is strong recommendation strength in favour of a DOG-appropriate behaviour, but the behaviour is punished because it is incorrect. This strong recommendation followed by negative reward triggers expansion of receptive fields in modules with strong internal activity, even if the number of active modules already reaches the minimum required level. There is a reasonable probability that these modules will not detect their receptive fields in a repetition of the DOG instance, making it possible to discriminate between the instances in the future.
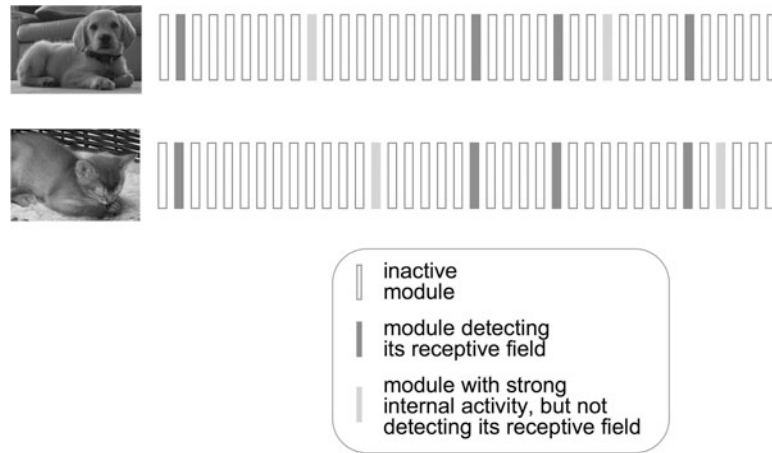

Figure 6 Managing discrimination problems. Initially, the visual similarities between the two pictures result in activation of the same sets of columns at the level of receptive field complexity that discriminates between object categories. Hence the same behaviour will be recommended in response to the two pictures (such as saying "cat") even though one is a cat and the other is a dog. The contradictory consequence feedback in response to the same behaviour following activation of the same columns results in forced receptive field expansions which lead to different column sets in response to the two objects

Because receptive fields expand over time, there is a requirement to limit the total number that is activated in response to one input state. This limitation is achieved by inhibition between receptive fields so that only the most strongly present are detected. The means by which this is achieved will be discussed after the information model for a pyramidal neuron has been described (figure 9).

5.2.2 Information Model for a Cortical Column

The information model for a cortical column is definition and detection of a receptive field, including identification of the circumstances in which expansion of its receptive field could be appropriate. This information model can be understood by consideration of

figure 7. The cortical column information model also includes indirect activation of a receptive field under special circumstances in which the field is not actually present in the current sensory input state. Such indirect activations are behaviours that must be adequately recommended by currently active receptive fields.
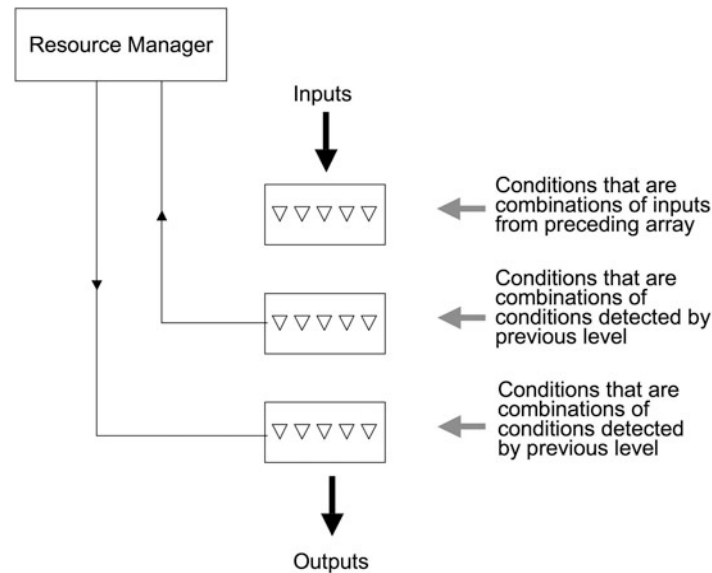

Figure 7 Information model for a cortical column. A simple cortical column model made up of three layers of principal neurons is illustrated. The conditions detected by neurons in the top layer are combinations of receptive field detections in the areas that provide inputs to the column. The conditions detected by neurons in the second layer are combinations of receptive field detections by neurons in the top layer. The conditions detected by neurons in the bottom layer are combinations of receptive field detections by neurons in the middle layer. Outputs from the bottom layer are the column receptive field detections and go to other cortical areas and also to subcortical structures where they are interpreted as behavioural recommendations. Receptive field complexities therefore increase somewhat from the top to the bottom layer. If in response to some input state there is no activity in the bottom layer but significant activity in the middle layer, this indicates that there is some similarity between the current state and other states that have contained the column receptive field. Such middle layer activity therefore indicated that if the receptive field of the column were expanded slightly, it would be detected in the current state. The strength of middle layer activity therefore indicates the degree to which receptive field expansion would be appropriate if required in response to the current input state. Outputs from the middle layer go to the resource manager, where they are used to determine the set of columns with the strongest degree of internal activity. If additional receptive field detections are required, then set of columns receive inputs from the resource manager that drive receptive field expansion.

In figure 7, there are three layers of principal neurons (generally pyramidal). The inputs that define the conditions detected in the top layer come from other areas of the cortex. The inputs that define the conditions detected in the middle layer come from the top layer, and the inputs that define conditions detected in the bottom layer come from the middle layer. The bottom layer provides outputs to other cortical areas and to subcortical structures where they are interpreted as behavioural recommendations. Note that this is a conceptual model, there could be more layers for various functional reasons (Coward 2005), and connectivity need not necessarily be sequential from top to bottom layer. From the point of view of the information model, the key point is that the conditions detected in the layer that provides outputs are less complex than the conditions

detected in earlier layers. A situation could therefore arise in which there in no output from the column as a whole, but significant condition detection in earlier layers. The implication of such a situation is that the current input state contains many of the conditions that contribute to the definition of the column receptive field but not sufficient to trigger receptive field detection. Hence the degree of activity of neurons in the middle layer is a good indicator of the degree to which expansion of the column receptive field would be appropriate if required in response to the current input state.

The requirement is to determine if receptive field expansions are required (based on the overall degree of activity of columns across the area) and if so to determine the columns with the highest degree of internal activity and trigger receptive field expansions in those columns. Although in principle this requirement could be met by all-to-all connectivity between columns, it is more efficient and effective to use a resource manager (Coward 1990, 2005, 2009a). In figure 7, the output from the middle layer of the column to the resource manager carries information on the internal activity of the layer, and the inputs from the resource manager drives receptive field expansions as described below.

5.2.3 Information Model for a Pyramidal Neuron
The information model for a pyramidal neuron is definition and detection of a receptive field within a column. There are a number of aspects to this information model. First, separate information conditions that constitute the receptive field must be defined by groups of elements, with a condition being detected if enough of the elements are in the appropriate state. Second, the receptive field of the neuron must be detected and an output generated if a significant proportion of the conditions are detected. Third, it must be possible to expand the receptive field in appropriate circumstances by adding appropriate conditions. Fourth, it must be possible to limit the total activity across an array of columns, ultimately by limiting neuron activity in an appropriate fashion. Fifth, it must be possible to activate the neuron in some circumstances in the absence of its receptive field.

The pyramidal neuron information model is conceptually illustrated in figure 8. In this information model, conditions correspond with groups of synapses on one branch of the apical dendrite. These synapses are connection points for axons coming from other pyramidal neurons. An incoming action potential causes a synapse to inject voltage potential into its branch. The magnitude of the injected potential is proportional to the weight of the synapse. In the conceptual figure, synaptic weights are in arbitrary units (w). As described in the next section, an injected potential increases rapidly after the arrival of the action potential and then decays more slowly. A significant proportion of these synapses must receive incoming action potentials for the total potential in the branch to reach a threshold enabling a potential contribution deeper into the dendrite, and thus raise the chance of the soma producing an output action potential indicating detection of its receptive field. Because of the decay in potentials, the inputs to one branch must occur within a relatively short time so that the potentials resulting from action potentials at different synapses add to each other. Synapses on one branch thus correspond with the elements defining one information condition, which is detected if a high proportion of the inputs are active.
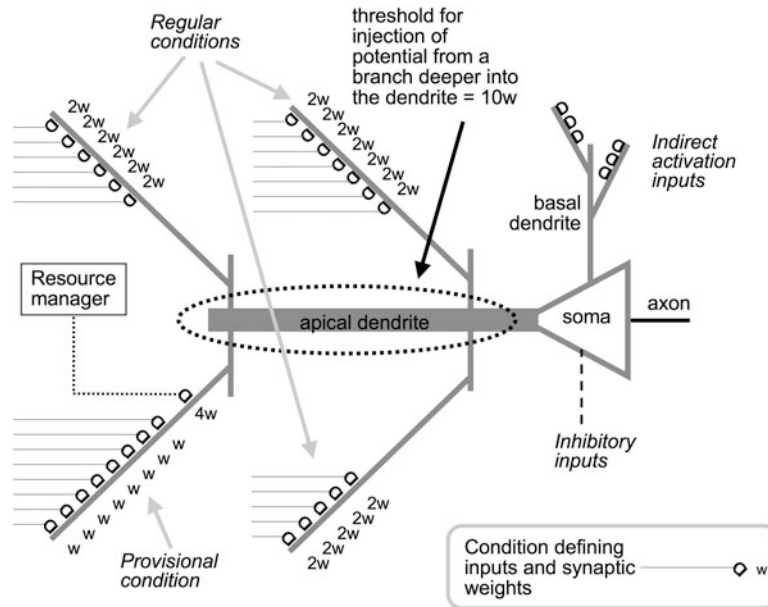
Figure 8 Information model for a pyramidal neuron. The neuron has a body (or soma) and two sets of dendritic trees. Inputs that define the conditions that in turn define the receptive field of the neuron synapse on different branches of the apical dendrite. If an action potential arrives at a synapse, it injects a voltage potential into the branch that is proportional to the weight of the synapse. If the total weight injected by action potentials arriving at all the synapses on a branch exceeds a threshold, the branch injects potential deeper into the dendrite. The group of synapses on a branch thus define an information condition. If the total potential injected into the dendrite reaches a high enough level (i.e. enough conditions are detected), the dendrite will inject potential into the soma. If the potential injected into the soma exceeds a threshold, the soma will generate an output action potential. The group of conditions defined on apical branches thus define the receptive field of the neuron. Inhibitory signals from local interneurons target the soma (or the proximal dendrites) and reduce the chance of the neuron producing an output. The basal dendrite can also inject potential to trigger an output action potential; the inputs to the basal dendrite come from sources that can generate such an output in the absence of the neuron receptive field

The neuron will only produce an output if a number of branches inject potential deeper into the dendrite within a relatively short period of time. The group of conditions corresponding with branches thus define the receptive field of the neuron, which is detected if a significant proportion of the conditions are detected.

A neuron can expand its receptive field by addition of new conditions. Such addition is achieved by means of provisional conditions such as the one illustrated in the lower left of figure 8. The total weight of all the synapses that define the elements of the condition is not sufficient for the branch to reach the threshold for potential injection deeper into the dendrite. However, there are additional synapses on the branch that come from the resource manager. If within a relatively short period of time, a high proportion of the condition defining elements are active and the inputs from the resource manager are also active, there will be enough potential in the branch to inject potential deeper into the dendrite. If a number of regular conditions are also being detected, then there will be enough potential to trigger firing of the neuron. If the neuron fires, a backpropagating acting potential into the branch increases the weights of any synapses that have recently received an action potential from their source neuron. This is the long-term potentiationmechanism observed by Bi and Poo (1998). The effect is that the weights of

the condition defining inputs are increased to the point that they can contribute to the firing of the neuron independent of the state of the inputs from the resource manager. Effectively, a new condition has been recorded.

A key requirement is that new conditions on a neuron must be similar to previously recorded conditions on the neuron or on other neurons in the same layer of the same column. As discussed earlier, similarity means that the elements defining the condition are the same as and/or often active in the past at the same time as the elements defining other conditions. There are a number of factors that ensure this similarity. First, for a new condition to be recorded, conditions already programmed on the same neuron must be detected. Second, the inputs to a provisional condition are selected from those available in the neighbourhood. Third, the inputs to a provisional condition are selected from those often active in the past when the neuron is also active. This final factor is achieved utilizing REM sleep (Coward 1990 etc.). In REM sleep, there is a rerun of a selection of past experience. Provisional connectivity is established between neurons often active at the same time during this rerun.

Limiting activity within a column and also across an array of columns is achieved using inhibitory interneurons with the connectivity illustrated in figure 9.

As described earlier, there is behavioural value in the ability to indirectly activate columns on the basis of different types of temporally correlated past activity. Inputs to indirectly activate neurons in a column must be segregated from inputs defining receptive fields. Hence these inputs target a different dendrite system, the basal dendrites.

The staged integration model described in this section is generally consistent with the ideas of Hausser and Mel (2003). However, the details of the staging could be different; the key requirement is to achieve appropriate receptive field expansions.
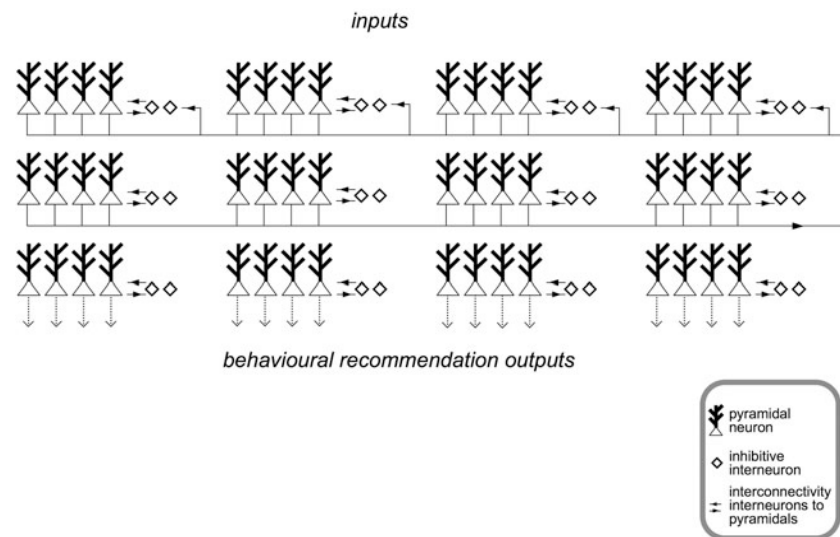

Figure 9 Interneuron connectivity to limit column activity. Interneuron outputs inhibit their targets. Interneurons target pyramidal neurons in the same layer of the same column. Interneurons can receive inputs from pyramidals in the same layer of the same column or from the same layer of columns in the same area. Connectivity between columns can be economized by biasing it in favour of connectivity between interneurons and pyramidals often active at the same time. Connectivity within a column limits overall activity within the column. Connectivity between columns effectively results in a competition between columns, reducing the total number of active columns.

5.2.4 Pyramidal Neuron Dynamics

As described earlier, a number of frequencies can be observed in the EEG which reflects the firing patterns of pyramidal neurons. In the recommendation architecture model, these frequencies reflect different information processing functions by networks of pyramidal neurons. In general terms, one action potential generated by a neuron indicates a detection of its receptive field at one point in time. The degree to which the receptive field is present is reflected in the rate at which action potentials are generated, and such rates averaged across many neurons result in the beta band frequencies. As discussed earlier, effective use of cortical resources requires the ability to maintain separate but simultaneously active populations of receptive field detections within the same resources. The gamma band reflects an information process that maintains a separation between simultaneous receptive field detections within different sensory objects in the same neuron network. The theta band reflects an information process that maintains a separation between groups of receptive fields that are detected at a sequence of different points in time.
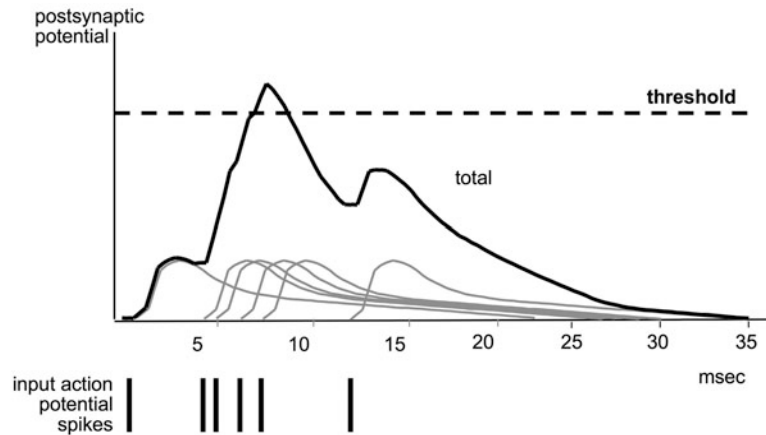

Figure 10 Leaky integration and integration windows in a pyramidal neuron. For simplicity, staged integration is omitted. Action potential spikes arriving at different synapses each result in the injection of potential into the neuron. This injected potential rises over a period of a couple of milliseconds, then decays again with a decay constant of the order of 10 ms. It is assumed that the synaptic strengths are all the same. The postsynaptic potential resulting from different spikes adds linearly, and the total potential must exceed a threshold to result in (or contribute to) firing of the neuron. Because of the decay constant, two spikes must arrive close together in time to reinforce each other. Unless a number of spikes arrive within a fairly short period of time (of the order of the decay constant), the threshold will not be reached. This situation thus supports the concept of an integration window or period within which a minimum number of spikes must arrive for there to be a contribution to target neuron firing.

The information model for the response of a pyramidal neuron to a sequence of action potentials is illustrated in figure 10. For simplicity, dendritic structure and staged integration is omitted for the figure, and it is assumed that the synaptic strengths are all the same. The potential injected by one action potential rises relatively rapidly after the arrival of the input action potential, peaks after 2–3 ms and then decays with a time constant of the order of 8 ms. This time constant implies that action potentials arriving significantly more than 8ms apart do not reinforce each other. Hence if the threshold for a

group of synapses to contribute to the firing of the neuron is several times the maximum potential injected by one spike, at least a minimum number of spikes must arrive within a period of the order of the time constant for the group of spikes to contribute to the firing of the neuron. The concept of an integration window is therefore useful, where at least a minimum number of spikes must arrive within the integration window to contribute to neuron firing. The integration window is not completely fixed, it depends both on the ratio of leaky integration peak to threshold and on how close together the spikes arrive, but simulations demonstrate that it is a viable concept (Coward 2004).

An important concept that makes the integration window functionally useful is frequency modulation as illustrated in figure 11. A frequency modulation signal applied to a neuron shifts each of its output spikes towards the nearest peak in the frequency modulation signal. Such a frequency modulation signal could be imposed by a regular sequence of spikes at the modulation frequency. Interneuron inputs generally inhibit the firing of their targets, but if an inhibitory spike arrives more than about 5 milliseconds before an excitatory spike, it adds to the excitatory effect (Gulledge and Stewart 2003). A regular stream of interneuron spikes will therefore tend to impose a frequency modulation on the outputs of pyramidal neurons.

One functional value of frequency modulation can be understood from figure 12. If unmodulated inputs from some visual domain are provided to an array of columns, then the degree of receptive field detection will be much lower than if the inputs are modulated. Hence a domain in the visual field can be selected by placing a frequency modulation on all the inputs from that domain, and the effect will be preferential receptive field detection (and therefore behavioural recommendation generation) with respect to that domain. This is the information model for the attention function. Note that if the modulation is placed on inputs, the outputs of the neurons will also be frequency modulated, and the modulation will propagate through following layers of neurons.
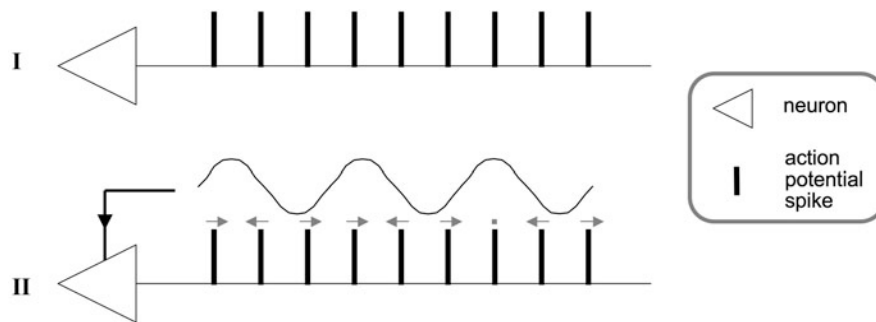

Figure 11 Concept of frequency modulation. In I, a neuron produces an output made up of a sequence of action potential spikes. For ease of explanation, the illustrated sequence is regular; the actual sequences produced by neurons are irregular but the concept applies in the same way. In II, a frequency modulation signal is applied to the neuron. The effect of the signal is to shift output spikes towards the nearest peak in the modulation signal
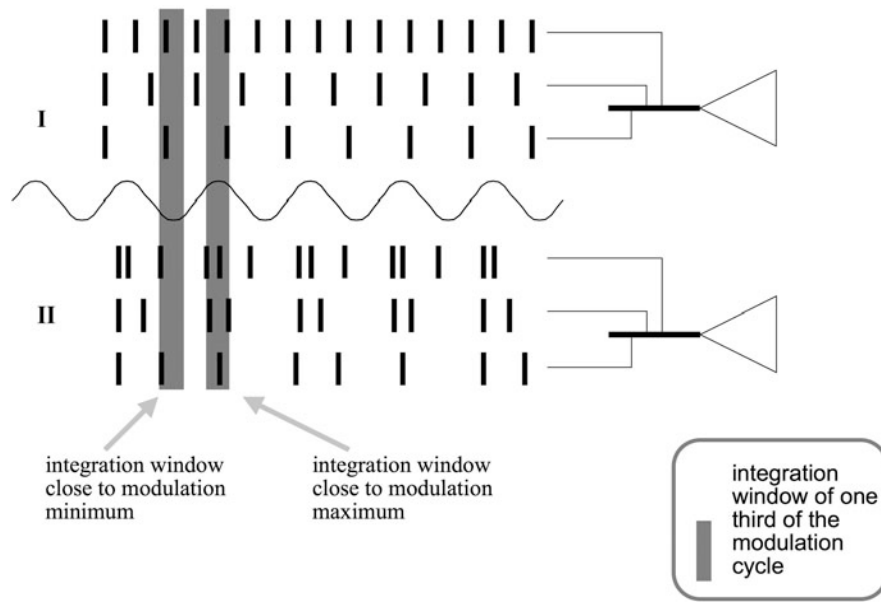
Figure 12 Frequency modulation of inputs from a source. In I, three unmodulated input sources target a neuron. Within any integration window, there is only a total of about two spikes. If at least 4 spikes must be present to generate an output, there will be no outputs. In II, the inputs are modulated. This results in the same or fewer spikes in the integration window around the modulation minimum, but more around the maximum. There are sufficient around the maximum to fire the target neuron. Hence if a group of inputs from one source (such as a domain in the visual field) are frequency modulated, receptive fields will be detected only within the inputs from that domain
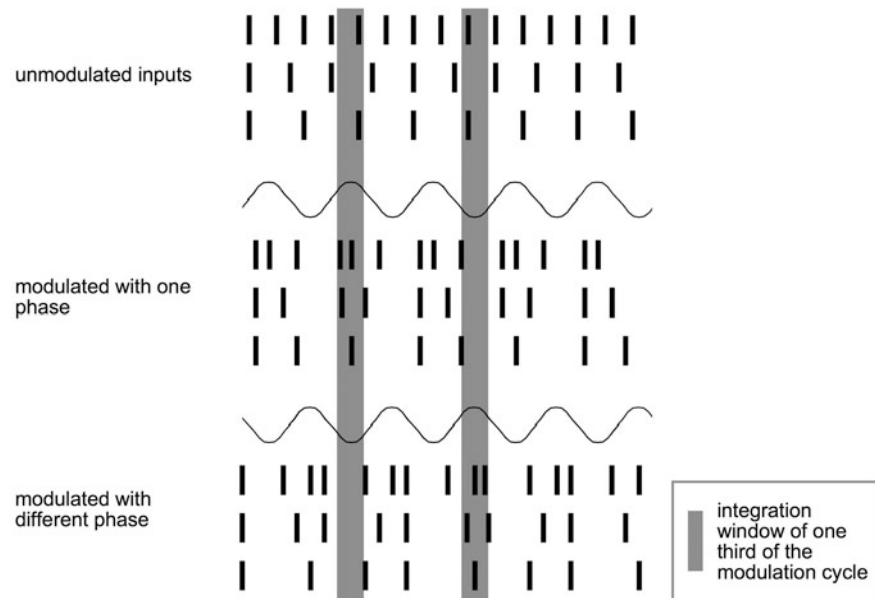

Figure 13 Frequency modulation enabling simultaneous receptive field detections in multiple objects without interference. Inputs from different sources are given different phases of the same frequency modulation. As a result, spikes from the different sources tend to arrive in different integration windows, meaning that receptive fields present within different input sources can be detected without interference in the same neural resources, even in the same neuron.

Another functional value can be seen from figure 13. If the interval between peaks in the modulation frequency is much larger than the integration window, then several integration windows can fit within one modulation cycle. If inputs from several different sources are modulated with the same frequency but at different phases, the spikes from the different sources will tend to arrive in different integration windows. Hence neuron receptive fields will be detected separately and independently within the different input sources. In other words, receptive fields within several different input sources (such as visual objects) can be detected simultaneously without crosstalk in the same neural resources.

The principle of maintaining a separation between different populations of receptive field detections can be extended to make use of multiple different frequencies. For example, in order to separate detection groups of receptive fields corresponding with sensory circumstances at a sequence of different points in time, a different frequency could be used.
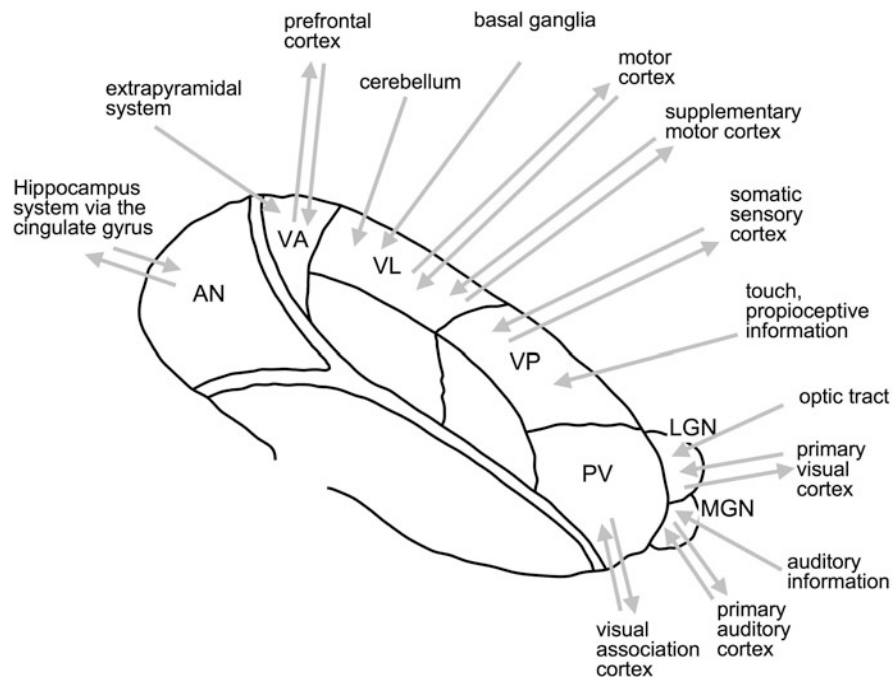

Figure 14 Organization of the thalamus into nuclei with reciprocal connectivity between a nucleus and specific cortical areas. In addition, different additional cortical areas provide inputs to each nucleus but do not receive inputs in return. All connectivity between the thalamus and the cortex are excitatory

5.3 Structure of the Basal Ganglia and Thalamus

The basal ganglia and thalamus are organized into nuclei: clusters of neurons separated by regions containing mainly axons. The major nuclei and connectivity of the basal ganglia and thalamus can be understood by consideration of figures 14–17.

The thalamus is made up of a number of major nuclei as illustrated in figure 14. Each nucleus has strong reciprocal excitatory connectivity with a different cortical area, plus inputs from some other areas. Between each nucleus and the cortex, there is a nucleus called the thalamic reticular nucleus (TRN). As illustrated in figure 15, all axons between

a thalamic nucleus and its associated cortices pass through a sector of the TRN, which regulates the information flow (Guillery et al. 1998).

There are a number of separate nuclei that make up the basal ganglia. These nuclei and the major connectivity between them and with the thalamus and cortex are illustrated in figure 16.

There are in fact several parallel paths from the cortex through the striatum, GPi/SNr and thalamus and back to the cortex (Alexander et al. 1986). These paths start and end in a cortical area different for each path, and each path goes through different subnuclei of the striatum, GPi, SNr and thalamus. In each path, the striatum also receives cortical inputs from some additional areas, but there is no return connectivity from the path to those areas.
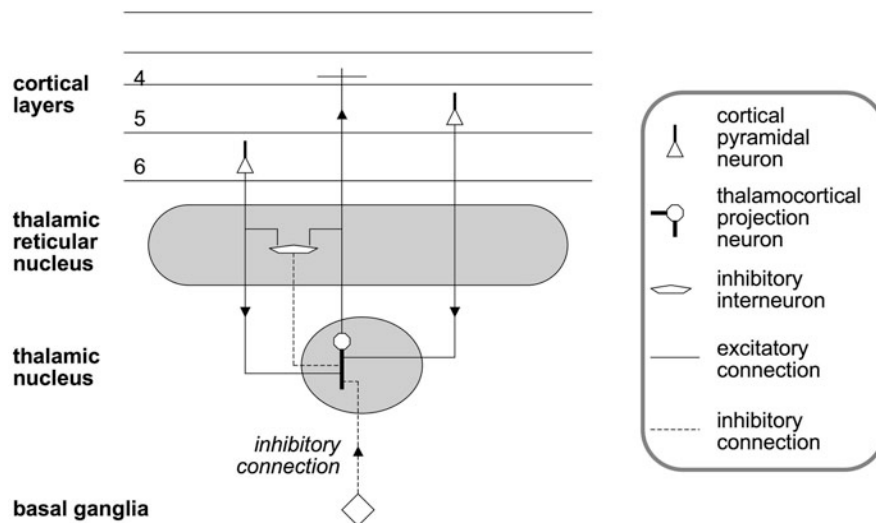

Figure 15 All connectivity between a thalamic nucleus and the cortex passes through a sector of the TRN. The primary connection paths for one type of thalamic nucleus are illustrated (Guillery et al. 1998). The principal neurons of the thalamic nucleus, thalamocortical projection neurons, send excitatory connections to layer 4 of the cortex. Thalamocortical neurons receive excitatory inputs from pyramidal neurons in layers 5 and 6. Axons from layer 6 pyramidals and from thalamocortical neurons branch in the TRN to form synapses on the inhibitory interneurons that are the most common cells in the TRN, but the axons from layer 5 do not form such synapses. These inhibitory interneurons project to thalamocortical neurons. In addition, thalamocortical neurons receive inhibitory inputs from neurons in the basal ganglia

Individual neurons in the striatum and in the globus pallidus correspond with specific aspects motor activity such as the direction of arm movement but not to the underlying pattern of muscle activity associated with that movement (Crutcher and DeLong 1984; Mitchell et al. 1987).

The nucleus accumbens is sometimes regarded as part of the basal ganglia. As illustrated in figure 18, it gets substantial inputs from the amygdala and from the orbitofrontal cortex, and its outputs target the GPi and SNr nuclei of the basal ganglia.

5.4 Information Models for the Thalamus and Basal Ganglia
The general information model for these structures is interpretation of cortical receptive field detections as behavioural recommendations and determining and implementing the most strongly recommended behaviour. Reward feedback following a behaviour adjusts

recently utilized recommendation weights. There are five components to this information model. The first is determination of raw total recommendation weights for each behaviour. The second is a competition between different behaviours to determine the most appropriate. The third is modulation to ensure that one and only one behaviour is selected. The fourth is implementation of the one selected behaviour. The fifth is utilizing reward feedback to modulate the probability of the same behaviour being selected in similar circumstances in the future.
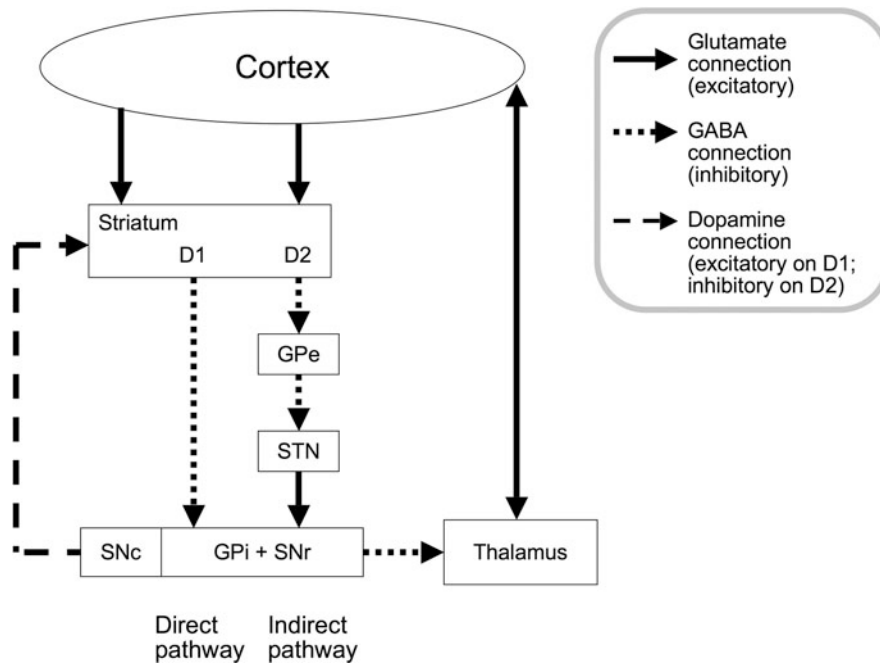

Figure 16 The major nuclei and connectivity of the basal ganglia. Excitatory inputs from the cortex arrive at spiny projection neurons in the striatum. Each such neuron receives one or two inputs from a very large number of different cortical pyramidals. Striatal spiny projection neurons inhibit their targets, and there are two populations of these neurons. D1 population neurons project directly to the globus pallidus internal segment (GPi) and substantia nigra pars reticula (SNr) and are excited by dopaminergic inputs from the substantia nigra pars compacta (SNc). D2 neurons project indirectly to the GPi and SNr via two additional nuclei, the globus pallidus external segment (GPe) and the subthalamic nucleus (STN). As a result, the D2 population neurons ultimately excite the GPi and SNr. D2 population neurons are inhibited by dopaminergic inputs from the SNc. The GPi and SNr generate constant (tonal) inhibitory outputs to the thalamus. The direct path therefore tends to reduce inhibition of the thalamus while the indirect path tends to maintain this inhibition. As described earlier, the thalamus has reciprocal excitatory connectivity with the cortex. There is no direct return path from the basal ganglia to the cortex

4.1 Information Model for the Thalamus

The thalamus implements selected behaviours. Such an implementation is generally the release of information into the cortex, between cortical areas, or out of the cortex. Principal cells in the thalamus excite different groups of cortical columns and correspond with the releasing of information from those groups. Thalamic principal cells are excited by a range of columns in their groups and elsewhere. This connectivity can be viewed as the cortical columns recommending the release of outputs from the group. However, the basal ganglia tonically inhibits the thalamic cells, and release will not occur unless the tonic inhibition is reduced. The release of information includes imposing frequency

modulation with an appropriate phase. This frequency modulation could be the role of the TRN, with TRN interneurons inhibiting thalamocortical neuron activity out of phase with the targetted modulation peaks.
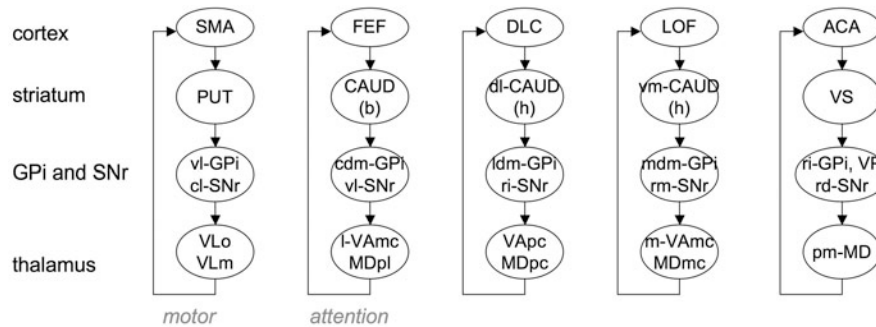

Figure 17 Parallel paths linking basal ganglia, thalamus and cortex. In each path, one striatal subnucleus receives inputs from one primary area, sends outputs to one pair of GPi and SNr subnuclei, which send outputs to one thalamic nucleus that connects reciprocally with the primary cortical area. In each path, a different group of cortical areas also provide inputs to the striatal subnucleus but does connect reciprocally with the corresponding thalamic nucleus. Abbreviations follow Alexander et al. (1986)
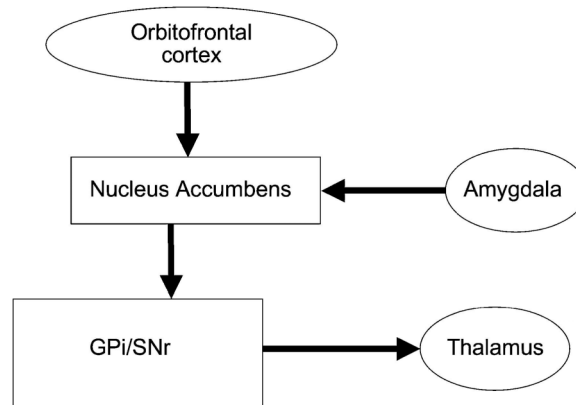

Figure 18 The primary connectivity of the nucleus accumbens. Following a behaviour, both the orbitofrontal cortex and amygdala provide information indicating the detection of circumstances in which reward feedback is appropriate. The nucleus accumbens interprets these inputs as recommendations in favour of adjustments of the weights in the GPi and SNr that resulted in the recent behaviour and applies the most strongly recommended adjustments. The weight adjustments therefore affect the probability of acceptance of similar behaviours in similar circumstances in the future

4.2 Information Model for the Striatum

The information model for the striatum is that different striatal projection neurons correspond with different behaviours or types of behaviour. Each neuron receives inputs from many cortical pyramidals, which can be interpreted as raw recommendation weights in favour of the behaviour corresponding with the striatal neuron. For reasons discussed below, there is a pair of striatal projection neurons corresponding with each behaviour, with similar inputs, but one population D1 and the other population D2.

Competition within the striatum is implemented by the extensive local axon collaterals of striatal projection neurons, which mainly target other striatal projection neurons (Somogyi et al. 1981). In addition, there are populations of striatal interneurons

that receive inputs from the cortex and target striatal projection neurons (Tepper and Bolam 2004).

5.4.3 Information Model for the GPi and SNr

The information model for the GPi and SNr is that individual principal neurons correspond with behaviours or types of behaviour. A principal neuron produces steady (tonic) inhibitory output to thalamic principal neurons corresponding with the release of cortical outputs that implement the behaviour corresponding with that principal neuron. With this connectivity arrangement, if a principal neuron in the GPi and SNr is inhibited, the result will be implementation of its corresponding behaviour.

A GPi and SNr principal neuron corresponding with a particular type of behaviour is targetted over the direct path (figure 16) by the D1 population striatal neuron corresponding with the same behaviour. This connectivity is inhibitory and therefore tends to encourage the implementation of the behaviour. The GPi and SNr neuron is targetted over the indirect path by D2 population neurons corresponding with different behaviours. This connectivity is excitatory and therefore tends to discourage the behaviour. GPi and SNr therefore supplement the competition between alternative currently recommended behaviours.

5.4.4 Information Model for the Nucleus Accumbens

The nucleus accumbens is observed to be associated with rewards (e.g. Ritz 1999). It targets the area of the basal ganglia in which the competition between different alternative behaviours occurs. The information model is therefore that following a positive reward it increases the weights that favoured recently accepted behaviours and decreases the weights that opposed them and vice versa for a negative reward.

5.5 Structure of the Hippocampal System

The hippocampus proper is made up of the CA fields and the dentate gyrus. The hippocampal system is made up of the hippocampus proper, some associated cortices and some associated subcortical structures. There is extensive connectivity between these structures as illustrated in figure 19.

All cortical areas except primary sensory project to the perirhinal or the parahippocampal cortices. These cortices project to the entorhinal cortex, which in turn projects to all the components of the hippocampus proper. CA1 generates outputs that go back through the entorhinal, perirhinal and parahippocampal cortices to the cortical areas from which inputs are derived.

The CA fields are cortex like, but with just one layer of pyramidal neurons. Within the hippocampus proper, there are two positive feedback loops: dentate gyrus granule cells excite mossy cells and CA3 pyramidals excite large numbers of other CA3 pyramidals. Each granule cells excites a small number of CA3 pyramidals and (via CA3 interneurons) inhibits a much larger number of CA3 pyramidals.
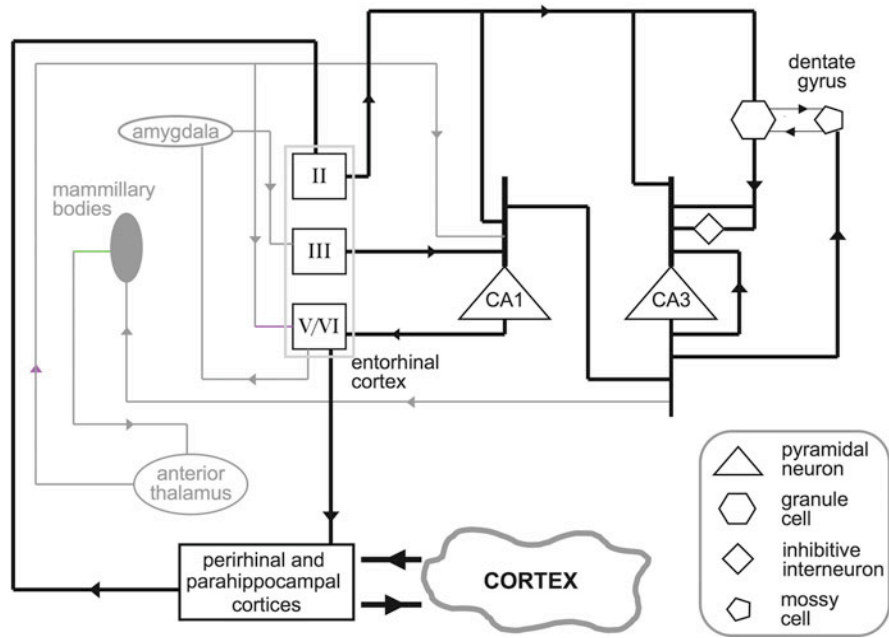
Figure 19 The structure of the hippocampal system. Outputs from all cortical areas except primary sensory areas go through the perirhinal and parahippocampal cortices to the entorhinal cortex and from there to all regions of the hippocampus proper. There is massive connectivity between the various CA fields and the dentate gyrus of the hippocampus proper. Outputs from CA1 return via the subicular complex, and the entorhinal, perirhinal and parahippocampal cortices to the cortex. In addition, there are several connectivity loops involving subcortical structures. There is connectivity from the CA fields to the mammillary bodies, connectivity from the mammillary bodies to the anterior thalamus and so back to CA1. The subicular complex and entorhinal cortex project to the amygdala, which it turn projects back to those same structures

5.6 Information Model for the Hippocampal System

The information model for the hippocampal system is definition of which cortical columns will expand their receptive fields at each point in time. Identification of the most appropriate columns for such expansion uses two sources of information. One is information on the current internal activity of all cortical columns. The other is information on groups of columns that have tended to expand their receptive fields at similar times in the past.

The information model for the associated cortices is illustrated in figure 20. Columns in the PHC and PRC have receptive fields that are internal activity by groups of cortical columns that have tended to expand their receptive fields at similar times in the past, where internal activity is indicated by outputs from column middle layers. Bottom layer outputs from a PHC or PRC column target the cortical columns that make up their receptive fields. These bottom layer outputs are not generated unless there is both receptive field detection and input from the ERC. Such an input from the ERC triggers a PHC or PRC column output, with receptive field expansion of the PHC or PRC column if required. Such a receptive field expansion would therefore add a new group of cortical columns about to expand their receptive fields at the same time.

Columns in the ERC have receptive fields that are internal activity by groups of PHC and PRC columns that have tended to expand their receptive fields at similar times in the past. Bottom layer outputs from an ERC column target the PHC and PRC columns that make up their receptive fields. Such bottom layer outputs are only generated if there is

both receptive field detection and inputs from CA1 pyramidals. Inputs from CA1 pyramidals trigger bottom layer outputs with receptive field expansions as required. Again, such receptive field expansions add groups of PHC and PRC columns about to expand their receptive fields at the same time.
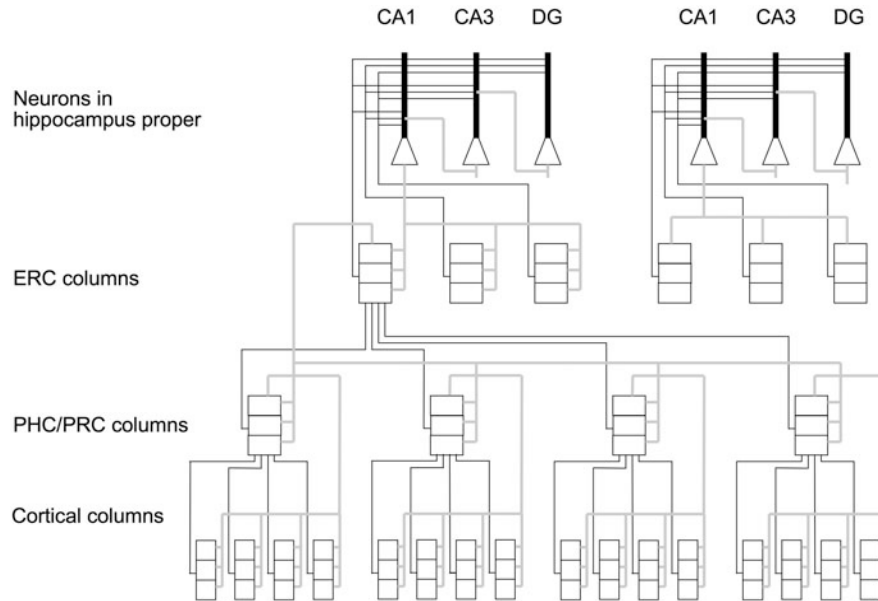

Figure 20 Information model for the hippocampal cortices. Columns in the PHC and PRC have receptive fields which are internal activity of groups of cortical columns that have tended to expand their receptive fields at the same time in the past. One cortical column may appear in the receptive fields of a number of PHC or PRC columns. An output from the bottom layer of one of these PHC and PRC columns encourages receptive field expansion in the cortical columns that form its receptive field. However, to produce a bottom layer output, a PHC or PRC column requires both strong receptive field detection and an input from ERC columns. Such an ERC input encourages PHC or PRC output, with receptive field expansion if necessary. Similarly, columns in the ERC have receptive fields which are internal activity of groups of PHC and PRC columns that have tended to expand their receptive fields at the same time in the past and require both strong receptive field detection and input from CA1 pyramidals to generate an output back to the PHC and PRC columns that form their receptive fields. CA1 pyramidals, CA3 pyramidals and granule cells all have receptive fields which are internal activity of groups of EC columns that have tended to expand their receptive fields at the same time in the past

The connectivity within the hippocampus proper is illustrated in figure 21. CA1 pyramidals and CA3 pyramidals all have receptive fields which are groups of ERC columns that tend to have expanded their receptive fields at the same time in the past. Granule cells have similar receptive fields, but poorly focussed so that they correspond more with groups of ERC columns that tend to be active at similar times in the past. The degree of input activity from the ERC indicates the degree of familiarity in the current sensory input state to the cortex. If there is a high level of such input activity across all cortical areas, there is a high degree of granule cell activity. This high level of granule cell activity means that inhibition of CA3 pyramidals is predominant and there is no CA3 activity and therefore no CA1 activity and no cortical receptive field expansions. If there is a degree of novelty in the current sensory input state to the cortex, some of the ERC input will be lower. Lower granule cell activity allows development of some CA3

pyramidal activity. The internal feedback within CA3 means that activity tends to develop in a group of CA3 pyramidals corresponding with large groups of cortical columns that have strong internal activity and also have tended to expand their receptive fields at the same time in the past. Direct inputs from granule cells encourage CA3 pyramidal receptive field expansion if required.
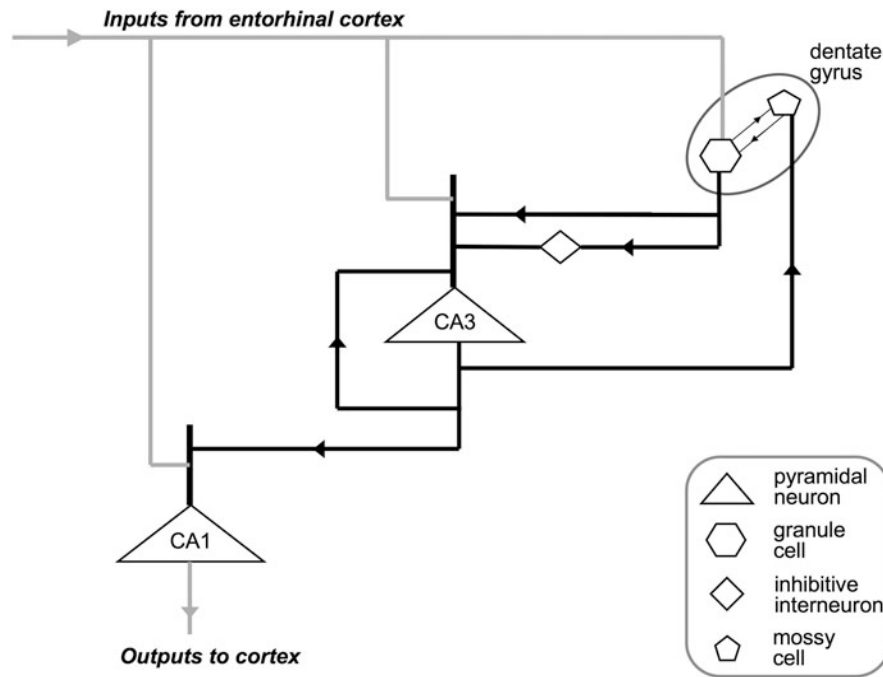

Figure 21 Connectivity within the hippocampus proper. CA1 pyramidals, CA3 pyramidals and dentate gyrus granule cells all have receptive fields that are groups of ERC columns that tend to expand their receptive fields at the same time. However, the CA1 pyramidals have sharply focussed receptive fields, CA3 pyramidals less sharply focussed, and granule cells relatively poorly focussed. There are two positive feedback loops: Each CA3 pyramidal excites a large number of other CA3 pyramidals, granule cells excite mossy cells and mossy cells excite granule cells. The two positive feedback loops are linked: CA3 pyramidals exciting mossy cells; granule cells excite a small number of CA3 pyramidals and (via CA3 interneurons) inhibit a much wider range of CA3 pyramidals. CA3 pyramidals target CA1 pyramidals

The projections from CA3 pyramidals to mossy cells mean that as CA3 pyramidal activity increases, granule cell activity will increase. Increasing granule cell activity will eventually reach a level where CA3 interneuron activity will cut off further increases in CA3 pyramidal activity. The lower the level of input from the ERC to granule cells, the higher the eventual level of activity in CA3 pyramidals. In other words, the degree of CA3 pyramidal activity will be proportional to the degree of novelty in the current sensory input state to the cortex.

The outputs of CA3 pyramidals target CA1 pyramidals with similar receptive fields. These outputs encourage CA1 pyramidal outputs, with receptive field expansion if required. CA3 activity is also communicated to the mammillary bodies. The mammillary bodies indicate to the anterior thalamus when the feedback competition within CA3 has settled down, and the anterior thalamus encourages CA1 outputs. To produce an output, a

CA1 pyramidal therefore requires inputs from the ERC indicating receptive field detection, inputs from CA1 pyramidals indicating that a similar receptive field has won the competition for receptive field expansion and inputs from the anterior thalamus indicating that CA1 pyramidal activity has stabilized.

An output from a CA1 pyramidal triggers a cascade of activity and receptive field expansion through the ERC, the PHC and PRC to the cortical columns appropriate for receptive field expansion. These receptive field expansions bring the overall level of cortical activity up to the minimum level required to generate a wide enough range of behavioural recommendations to support a high integrity behaviour selection.

Receptive field expansions by granule cells are unguided. Hence they are poorly focussed on groups of columns that expanded their receptive fields at similar times in the past. However, these granule cells provide receptive field expansion guidance to CA3 pyramidals, which therefore have more sharply focussed receptive fields. CA1 pyramidals have receptive field expansions guided by CA3 pyramidals and therefore have the most sharply focussed receptive fields, appropriate for guiding the cortex.

The amygdala plays a role in emotion (Anderson and Phelps 1997; Whalen 1998). The information model for the role of the amygdala within the hippocampal system is that in emotional situations it increases the level of signals driving receptive field expansion in certain cortical areas, in particular those areas detecting receptive fields that discriminate between different general situations.

At a particular location, there will be a relatively constant surrounding sensory environment. As a result, the large group of cortical columns activated in response to this environment will have tended to expand their receptive fields at the same time in the past. Many of the receptive fields in the hippocampal system will therefore correspond with particular locations, leading to the existence of the observed place fields in the higher levels of the hippocampal system (Fyhn et al. 2004; Leutgeb et al. 2004).

In order to support indirect activation on the basis of receptive field expansion shortly after currently active columns, the simultaneous activity of hippocampal system neurons with receptive fields corresponding with different large groups of cortical columns that have expanded their receptive fields at a sequence of points in time will be necessary. For example, a sequence of place fields must be maintained active during physical movement. These different hippocampal receptive field detections must only interact in a controlled fashion, and their separation is managed by the theta frequency in the EEG.

This use of the theta frequency to prevent interaction between different neural populations has some similarities with the proposal of Hasselmo et al. (2002) that different functions (learning and retrieval of earlier learning) are supported by activity in different phases of the theta cycle. However, the recommendation architecture information model is somewhat more consistent with the observations that the phase of the theta modulation shifts in a regular fashion across place fields (O'Keefe and Recce 1993).

5.7 Structure of the Cerebellum

The cerebellum has an outer sheet of tissue surrounding a body of white matter (i.e. axons) and a core of nuclei. The sheet has three major layers. The innermost layer is made up of very large numbers of granule cells. The middle layer is made up of Purkinje

cells. The outer layer is made up of the dendritic trees of the purkinje cells, penetrated by the axons of the granule cells.

One major source of excitatory inputs to the cerebellum is inputs containing motor information derived from the motor cortex and the spinal cord, via a nucleus called the inferior olive. These inputs target both the cerebellar nuclei and Purkinje cells. One input from the inferior olive targets a small set of Purkinje cells, and one Purkinje cell receives inputs from only one inferior olive input. A second source of excitatory inputs to the cerebellum is inputs containing sensory information derived from the cortex, via a structure called the pons. These inputs target both the cerebellar nuclei and granule cells.

Purkinje cells receive very large numbers of excitatory inputs from granule cells, each input very weak. The single axon from the inferior olive that targets a Purkije cell makes multiple synapses and strongly excites the cell. Purkinje cells make inhibitory connections on to cells in the cerebellar nuclei.

Excitatory outputs from the cerebellar nuclei target both the spinal cord through the brain stem and the thalamus.
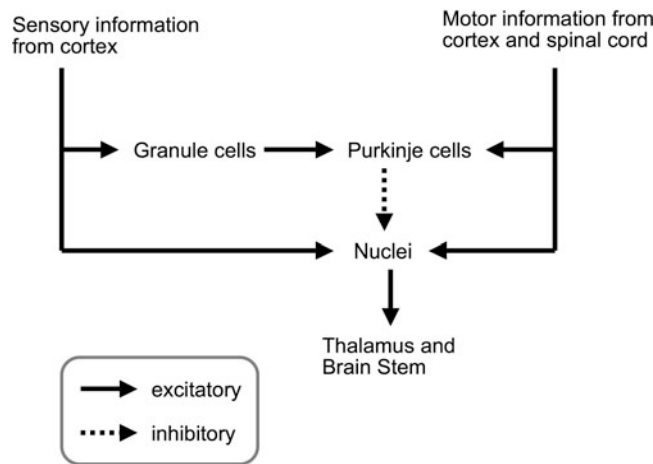

Figure 22 Information model for the cerebellum. A Purkinje cell corresponds with a "next behaviour". It receives an input containing motor information indicating that its preceding behaviour has taken place. It also receives inputs containing sensory information. Sensory information corresponding with completion of its "next behaviour" is weak compared with any other sensory input. Hence a Purkinje cell will be active from when its preceding behaviour is indicated until the completion of its own behaviour is indicated. Cerebellar nuclei cells correspond with different behaviours, and a Purkinje cell inhibits all except its corresponding next behaviour

5.8 Information Model for the Cerebellum

The information model for the cerebellum is rapid and accurate implementation of a frequently used sequence of behaviours. The way this information model is partitioned between different parts of the cerebellum is illustrated in figure 22. The next step in a sequence is instantiated by a cerebellar nucleus cell and by a Purkinje cell. The Purkinje cell receives an excitatory input from the inferior olive indicating that the behaviour preceding its corresponding behaviour has occurred and excitatory inputs indicating the presence of a wide range of sensory circumstances. However, the weights of inputs indicating the sensory circumstances that follow the performance of its corresponding behaviour are weak. The Purkinje cell inhibits all cerebellar nuclei cells except the one

corresponding with its own behaviour. Hence when the programmed preceding behaviour occurs and the sensory inputs are not consistent with completion of its programmed behaviour, the Purkinje inhibits all except its programmed behaviour. Once the sensory inputs are consistent with completion of its behaviour, the Purkinje cell outputs cease. At the start of learning, a Purkinje has inputs from a very wide range of sensory circumstances, but during learning the inputs that are present soon after it produces an output are weakened by the long-term depression (LTD) mechanism (Linden 2003).

Outputs from the cerebellar nuclei drive the next behaviour. For sequences of motor behaviours, these outputs can go directly to the spinal cord via the brain stem. For cognitive behaviours only involving releases of information within the cortex, outputs go to the thalamus.

Note that a behaviour sequence programmed in the cerebellum biases behaviour acceptance in favour of a series of types of behaviour, but the actual behaviour at each point will be the recommendation of the type that is most strongly recommended across the currently active cortical column population. At each point in the sequence, a different column populationwill be active as a result of the previous step in the sequence. Hence a sequence could proceed even with damage to the cerebellum because all the recommendation strengths will be in place, but it will be slower and more prone to errors. This is consistent with observations that cerebellar damage results in a generalized tendency towards poorer cognitive performance rather than severe specific deficits (Bracke-Tolkmitt et al. 1989).

6 Modelling of Memory and Learning Phenomena

There are a number of types of information recorded in different parts of the brain that are relevant to memory and learning. These information types are summarized in Table 1, along with the ways in which the information changes and the physiological instantiation of the information and the change mechanisms.

6.1 Receptive Fields Stability and Memory

In general the receptive field of a column can expand slightly with learning but does not contract or make major shifts. This relative stability of receptive fields and their sets of associated behavioural recommendations is the critical factor supporting many observed memory phenomena. If a receptive field is detected within the current sensory input, its associated behavioural recommendations have a high probability of being currently relevant. However, because of receptive field stability, some additional receptive fields that are not detected but can be identified also have a reasonable probability of having relevant behavioural recommendations.

For example, suppose that the receptive field of a column is not being detected and the column is therefore inactive, but the column was recently active at the same time as a number of the columns that are currently active. There is a reasonable probability that the recommendations associated with that column could be relevant. Similarly, if a column is inactive, but has often been active in the past at the same time as a number of the columns that are currently active, again its recommendations could be relevant. Finally, if a column is inactive, but it expanded its receptive field at the same time in the past as a number of currently active columns, its recommendations could also be relevant. Furthermore, the behaviourally relevant past activity of the inactive column could be at

the same time as just before or just after the past activity of currently active columns. There is therefore potential behavioural value in indirect activation of columns on the basis of various types of past activity that is temporally correlated activity with the past activity of currently active columns. Such indirectly activated columns could in turn indirectly activate further columns on the basis of past temporally correlated activity and so on. An indirect activation behaviour applied to an already active column would have the effect of prolonging its activity.

Table 1. Different types of information that support memory and learning in the brain. The major types are receptive fields in the cortex that record similarity circumstances, recommendation weights in the basal ganglia that associate receptive field detections with behaviours, and behaviour sequences in the cerebellum that associate behaviours with frequent next behaviours. Many of the memory relevant behaviours and next behaviours recorded in the cerebellum are indirect activation behaviours acting on cortical columns.

Information type	Changes with time	Physiological instantiation of information	Physiological instantiation of changes
Receptive fields			
Receptive field detecting presence of a type of sensory circumstance		Columns in sensory cortices	
Receptive field detecting activity of a group of cortical columns with frequent past temporally correlated activity	Slightly expands IF only a slight expansion is needed to result in current detection AND *either* the total number of receptive field detections is low *or* contradictory reward feedback has occurred in the past following activation of currently active group of columns	Columns in anterior temporal cortices	LTP mechanism operating on pyramidal neurons within the changing column to add conditions to neuron receptive field definition
Receptive field detecting activity of a group of cortical columns with past temporally correlated receptive field expansion		Columns in cortices associated with the hippocampal system	
Recommendation weights			
Indirect activation of receptive fields on the basis of past temporally correlated activation	Starts off high, initially (minutes) decays rapidly, then (hours and days) more slowly. Decays more slowly if repeated or if an indirect activation is followed by a reward	Weights of connections from columns in the anterior temporal cortex into the basal ganglia	
Indirect activation of receptive fields on the basis of past temporally correlated receptive field expansion	Starts off high, initially (hours and days) decays slowly, then (weeks and months) more slowly. Decays more slowly or increases if an indirect activation is followed by a reward	Weights of connections from columns in the cortices associated with the hippocampal system into the basal ganglia	Inputs from nucleus accumbens drive synaptic weight changes in the basal ganglia
Behaviour sequences			
Bias in favour of one behaviour type whenever a behaviour of another type has just been completed	Increases with repetition of sequence if followed by positive reward	Cerebellar purkinje neuron corresponding with the first behaviour that inhibits all except second behaviour	LTD mechanism reducing synaptic weights of sensory circumstances corresponding with completion of first behaviour

However, if such indirect activations took place without any limitations, there would be large numbers of columns indirectly activated in all situations. In some situations, only the directly activated columns are necessary to select appropriate behaviour and indirectly activated columns could reduce the appropriateness of the selection. Hence indirect activations must themselves be behaviours that are recommended by currently active columns, with competitions determining whether such recommendations are accepted.

6.2 Development and Evolution of Indirect Activation Recommendation Strengths

If two columns are active at the same time, there is a chance that an indirect activation strength would be useful in the future, but the highest probability of being useful is soon after the simultaneous activity. However, if the simultaneous activity occurs repeatedly, there is a higher chance that an indirect activation would be useful further in the future. If an indirect activation strength is utilized and is followed by a reward, there is an even higher chance of future value.

Hence if two columns are active at the same time, they will immediately acquire recommendation strengths in favour of activating each other in the future, but this strength will decay rapidly with time. If the simultaneous activity often occurs, the recommendation strength will decay more slowly. If the recommendation strength results in an indirect activation and the activation is followed by a positive reward, the weight will be increased and decay will be much slower.

If two columns expand their receptive fields at the same time, there is a much higher probability of future behavioural relevance than for simple simultaneous activity. The initial indirect activation recommendation strengths will therefore decay much more slowly. Reward feedback following use of a recommendation strength will again increase and stabilize the weight. Frequent use of a recommendation weight in favour of activation on the basis of simultaneous receptive field expansion would mean that the columns are often active at similar times, and indirect activation weight on that basis would develop which could even become stronger than the original weight based on receptive field expansion.

6.3 Semantic Memory

The primary information mechanism supporting semantic memory is indirect activation of cortical columns on the basis of frequent past simultaneous activity. This will be illustrated by the example of learning the meaning of the word "bird". Visual experiences of different birds result in receptive field detections in column arrays that discriminate between visual elements, between visual features and between visual objects (as in figure 4). Auditory experiences of the spoken word "bird" result in receptive field detections in column arrays that discriminate between phonemes and between words (as in figure 4). Because there are some similarities between different birds, there are some columns in the visual objects array in figure 4 that tend to detect their receptive fields relatively frequently in different bird instances. There will be less consistency in the visual features array and even less in the visual elements array. Similarly, there will be some columns in the word array that frequently tend to detect their receptive fields in different auditory experiences of the word "bird".
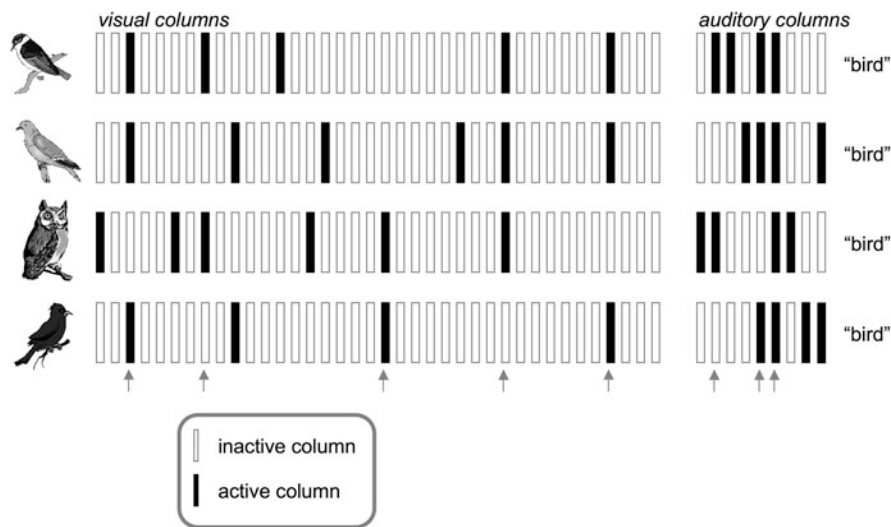
Figure 23 Information model for semantic memory. The column arrays that discriminate between visual objects and that discriminate between auditory words (as shown in figure 4) are illustrated for four situations. In each situation, a visual instance of a bird and auditory instance of the word "bird" are experienced simultaneously, resulting in activation of the columns that detect their receptive fields. Although the set of columns activated is different for each visual and each auditory experience, the similarity between the experiences means that some receptive fields tend to be detected relatively frequently in many visual bird instances and in many auditory "bird" instances, as indicated by the grey arrows. As a result, there is a set of visual columns that are often active at the same time as a set of auditory columns. The set of auditory columns therefore tend to acquire recommendation strengths in favour of indirect activation of the set of visual columns and vice versa. An instance of the word "bird" will contain a significant proportion of the auditory set, and there will therefore be significant total recommendation strength in favour of activating the visual set. The experience of activation of this visual set will be as if a visual bird instance that is an average of past experience were perceived, except that the columns close to visual input (i.e. arrays visual elements and visual features in figure 4) will not be activated, and the experience will therefore not be a visual hallucination.

In figure 23, the column activations in arrays visual objects and words are illustrated for a number of experiences a bird was perceived visually at the same time as the word "bird" was heard. In each visual experience, a different group of columns was activated, but because of the similarities between birds, there are some columns that tend to detect their receptive fields relatively frequently in different instances. Similarly, although a different group of auditory columns detect their receptive fields in each "bird" experience, some columns are active relatively frequently. There is therefore a set of visual object columns that are frequently active at the same time as a set of word columns. These two sets of columns will acquire recommendation strengths in favour of indirectly activating each other.

In any future experience of the word "bird", a significant subset of the auditory set will be activated, and this subset will have a strong total recommendation strength in favour of indirect activation of the visual set. This indirect activation will be experienced as if a bird that was an average over past experiences were experienced. Because there is less consistency in activation at the visual features and visual elements levels, the experience will not be a visual hallucination; it will be confined to higher receptive field

complexities. The activated visual object columns will have all their recommendation strengths, including relatively strong bird appropriate total recommendation strengths.

At a deeper level of description, the information model for this type of indirect activation is illustrated in figure 24. The indirect activation behaviour must win a competition for selection in the basal ganglia, and the behaviour must be implemented by the thalamus. Columns with receptive fields corresponding with groups of words auditory columns that are often active at the same time are required. These are the linking columns in figure 24. These linking columns target visual columns that are often active at the same time, but by connections on to basal dendrites of pyramidal neurons in those columns rather than on to the apical dendrites that define receptive fields.
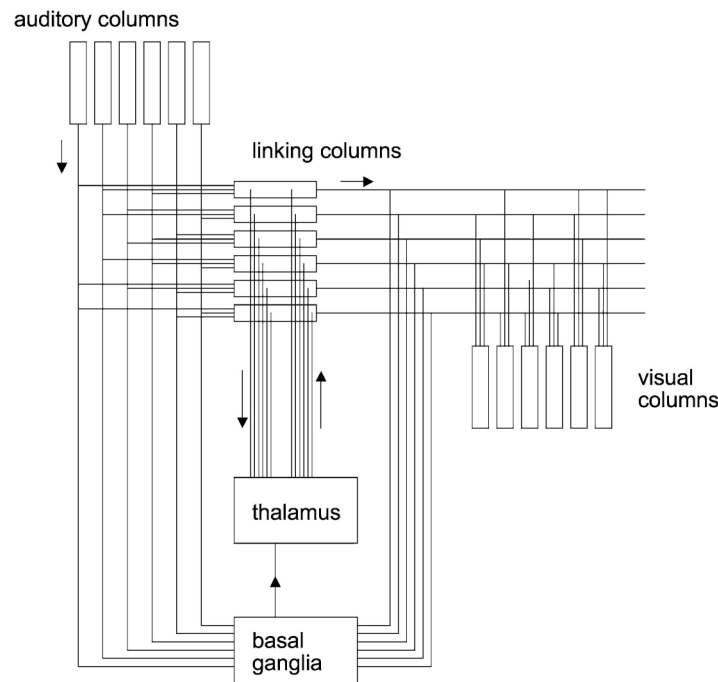

Figure 24 Semantic memory for word meanings supported by management of indirect activation on the basis of frequent past simultaneous activation. The linking columns have receptive fields corresponding with groups of auditory columns that are often active at the same time. The linking columns excite visual columns that are often active at the same time but not as elements in their regular receptive fields. Auditory columns and other columns including the linking columns target the basal ganglia components corresponding with the behaviour of indirect activation on the basis of frequent past simultaneous activity. These connections into these components have different synaptic weights which correspond with recommendation strengths in favour of the indirect activation behaviour. If the behaviour is accepted by the basal ganglia, the outputs from the linking columns are released by the thalamus to drive visual column activations. Note that the outputs from the linking columns recommending the indirect activation behaviour to the basal ganglia are from a different column layer to the outputs to the visual columns.

The receptive fields activated in the experience of a semantic memory will be those active during sensory or motor processing, and one receptive field will be present in instances of many different categories of object. Neuroimaging is consistent with this picture (e.g. Martin 2007). However, the activation of a semantic memory requires a cortical area in which columns have receptive fields corresponding with groups of columns elsewhere that are often active at the same time. Damage to this area would

result in general loss of semantic memory capability. Damage to the most anterior portions of the temporal cortices results in this type of general loss of semantic memory capabilities (Rogers et al. 2006).

6.4 Working Memory

In the recommendation architecture model, the working memory of an object is simple the activation of the columns corresponding with the semantic memory of the object. The primary information mechanism supporting working memory is frequency modulation of receptive field detections within different sources. The cognitive significance of the mechanism can be understood by reference to figure 25. Suppose there are three objects in the visual field, a dog, a cat and a tree, with the cat up the tree and the dog barking at the cat. The appropriate behaviour in response to the group (e.g. chasing the dog away) needs information about each of the objects that must be derived from receptive fields detected within those objects. For example, relevant information could include whether the cat is my pet, whether the dog looks fierce and the height of the tree. Receptive fields are detected within the three objects in the objects array in figure 25, but detections must be kept separate in this array (e.g. must not have information indicating a barking cat, or a fierce tree, etc.). However, the information derived from the three objects must be integrated in the groups array to activate columns with appropriate recommendation strengths.
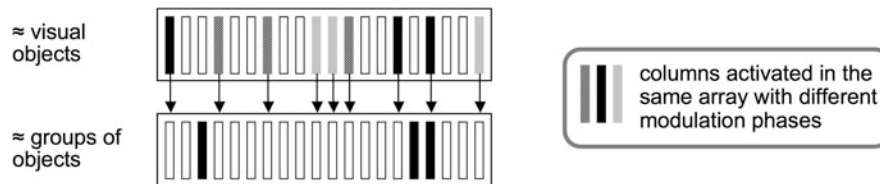

Figure 25 Cognitive significance of frequency modulation in maintaining independent populations of active cortical columns. Two arrays of visual columns are illustrated, one at a level of receptive field complexity effective for discriminating between objects, the other groups of objects. The receptive fields that discriminate between groups of objects need inputs from receptive fields that discriminate between objects. The modulation mechanism makes it possible for attention to be shifted to three different objects with column populations in response to the different objects being maintained active simultaneously but without interaction in the objects array until all the members of the group have been attended. The outputs from the three populations are then brought into the same phase and released to the groups array, where columns with recommendation strengths appropriate for responding to the group of objects then detect their receptive fields.

The frequency modulation mechanism makes it possible to place a different phase of modulation on the receptive fields detected within the dog, cat and tree. As a result, the receptive field detections are kept independent, even though they are all active in the same cortical array, potentially in some of the same columns or even some of the same pyramidal neurons. When all three populations are active, their outputs can be brought into the same phase and released to the groups array where receptive fields within the group are detected.

This type of mechanism supports the observed working memory limits on the number of different objects that can be retained in working memory at the same time. Cowan (2000) suggested a limit of approximately four items, based on performance discontinuities such as errorless performance in immediate recall when the number of

items is less than four and sharp increases in errors for larger numbers. In information model terms, the primary limit is defined by the number of active column populations which can be maintained independently. This number is approximately the ratio of the period of the modulation frequency to the postsynaptic potential decay constant. If the gamma frequency in the EEG is interpreted at the modulation signal on the basis of its association with attention (Muller et al. 2000), a 40Hz signal has a period of 25ms. With a postsynaptic decay constant of 10 ms (Fan et al. 2005), this would allow of the order of 2–3 independent column populations. The actual number would depend upon the gamma band frequency (which can vary from 20 to 50Hz) and the shape of the postsynaptic potential decay curve.

The content of an item in working memory is defined by a group of columnar receptive fields in a cortical area that are active at the same phase of frequency modulation. The column receptive fields are the same fields that record the information which makes up the content of semantic and episodic memories. This is consistent with the observations that there is considerable overlap in the cortical areas active during working memory and declarative memory tasks (Cabeza et al. 2002).

At a psychological level, the requirement is to keep information relevant to different objects active simultaneously, without interference. In particular, it is sometimes important to allow processing of a new sensory object while preserving information about an earlier object. However, it will not in general be necessary to keep information on the earlier object active at all levels of receptive field complexity, only at the levels with relevant behavioural recommendation strengths. In addition, there is a biological cost to maintaining information active, both the direct cost of the activation and the indirect cost of occupying resources that therefore cannot be used for other purposes. Hence maintaining information active in a specific cortical area must be a behaviour that will compete with other alternative behaviours for acceptance.

This competition for acceptance will depend upon what else is already active in different cortical areas. There is therefore a requirement for a cortical area with receptive fields corresponding with, for example, the number of different groups of columns in another area that are active with different phases of frequency modulation. Damage to this area would result in general deterioration of working memory capabilities. Neuroimaging indicates higher activity in Brodmann's area 40 in the left parietal cortex during working memory tasks (Cabeza et al. 2002), and a patient with damage to the left parietal lobe showed a deficit in working memory but not declarative memory (Warrington and Shallice 1969; Shallice andWarrington 1970).

If we are asked to imagine a number of different objects, keeping the images separate, the maximum number is 3–4, consistent with Cowan's (2000) limit. However, objects can be remembered as groups, and it is possible to remember objects as both visual and verbal terms. Cycling between these representations makes it possible to increase the apparent size of working memory, hence the "magic number" of seven (Miller 1956).

This model has some analogies with the proposal that different working memories are encoded as a different high frequency (40Hz) subcycle of a low frequency (10Hz) oscillation (Lisman and Idiart 1995). However, their model leads to a working memory content of seven, which is significantly higher that the actual experimental number.

6.5 Episodic Memory

The primary informationmechanism supporting episodicmemory is indirect activation of cortical columns on the basis of past simultaneous receptive field expansion. This will be illustrated by the example of recalling the news of the terrorist attack exploding a bomb in a bar in Bali. The state of column activation at some point while television news of the bombing was being viewed is shown in figure 26.

In the figure, attention is paid to a sequence of visual objects on the television screen. Different sets of cortical columns in the objects array detect their receptive fields within each object, and several sets corresponding with different objects are maintained active simultaneously at different phases of frequency modulation. Outputs from the different sets are then synchronized and released to the groups array where a set of columns detect receptive fields. The outputs from several sets of columns in the groups array are in due course synchronized and released to the groups of groups array.
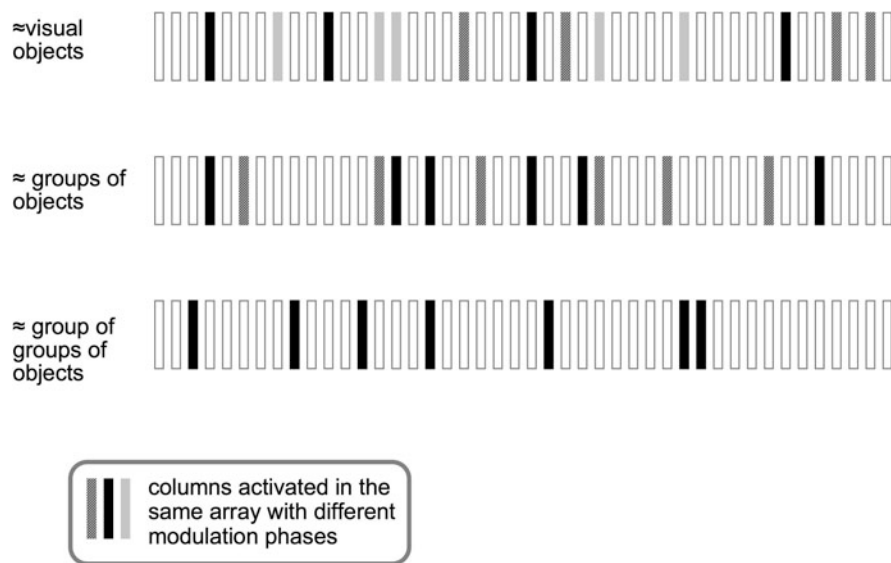

Figure 26 Conceptual representation of the cortical columns activated at some point during viewing of television news of the first Bali bombing. The pictured activation was built up by a series of steps. Several objects on the television screen were sequentially the focus of attention, resulting in separate activations in the objects array, at different phases of frequency modulation. Outputs from these separate activations were synchronized to the same phase and released to the groups array, an active column population was generated in that array at a specific phase of frequency modulation, and the populations in the objects array were extinguished. More visual objects generated activations in the objects array, their outputs were again synchronized, and a second independent population generated in the groups array at a different modulation phase from the first. In due course, the outputs from several independent populations in the groups array were synchronized to the same phase and released to the groups of groups array where a column population was activated.

There is relatively little novelty to the individual visual objects, which might include objects typical of Bali, of bars and of explosion aftermath, all of which are separately familiar. Hence little receptive field expansion is required in the objects array to achieve the minimum required number of column receptive field detections. However, at the groups level there will be rather more novelty, and hence a significant degree of receptive field expansion will be required, while at the groups of groups level there will be

significant novelty, and substantial numbers of receptive field expansions will be required.

Now suppose that later the words "Bali" and "bombing" are heard. By the semantic memory mechanism, sets of columns will be indirectly activated in the visual objects array corresponding with the visual experiences that have often occurred at the same time as the two words. The outputs from these sets would be synchronized and released to the groups array where they would generate a set of active columns. However, the overlap between the resultant column populations and the columns active in figure 26 will be slight.

If now the recommendations of the currently active columns in favour of indirect activation of other columns on the basis of past simultaneous receptive field expansion are accepted, a secondary population will be activated. This secondary population will contain columns that have expanded their receptive fields at the same time in the past as a significant subset of the columns activated in response to the words. A tertiary population could then be generated on the same basis from the secondary population. This sequence of indirect activations will tend to result in a population in which all columns tend to have expanded their receptive fields at the same time in the past. If the initial words are well chosen the final population will be a good approximation to the population active at the time of the original experience, but without the activations in the earlier cortical arrays (visual features and visual elements) in which the degree of receptive field expansion will be minimal. Nevertheless, this reconstructed population will have appropriate recommendation strengths in favour of describingwhat was seen. In other words, an episodic memory of the original experience has been constructed.

Column recommendation strengths in favour of activating other columns that expanded their receptive fields just before or just after expansion in the column provide the ability to move through the sequence of experiences to reconstruct an episode.

As discussed earlier, cortical columns in the parahippocampal, perirhinal and entorhinal cortices and pyramidal neurons in CA1 and CA3 define receptive fields corresponding with groups of cortical columns that have tended to expand their receptive fields at the same time. These receptive field definitions result from the primary resource management role of the hippocampal system. Hence these hippocampal system structures can be used to reconstruct episodic memories, with the restriction that outputs that drive receptive field expansions cannot be used without generating unnecessary and undesirable receptive field changes. The outputs to drive episodic memory recall must therefore come from different layers in the PHC, PRC and EC columns from those layers driving receptive field changes.

Note that if an episode is often recalled, there may be sufficient recommendation strengths acquired on the basis of frequent past simultaneous activity for that mechanism to take over from the receptive field expansion-based mechanism. This accounts for the observations that hippocampal damage has the largest effect on autobiographical memories and less effect on memories of personal information, notable public personalities and notable public events that may have been frequently recalled (Nadel and Moscovitch 1997).

This model of episodic memory recall involves a specific sequence of cortical behaviours: generation of activations in response to words; a series of indirect activations on the basis of past simultaneous receptive field expansion; evolution of the end

population on the basis of slightly later receptive field expansion; and generation of speech driven by the final population. This sequence must be learned and later performed rapidly and effectively. This sequence will therefore be instantiated in the cerebellum.

Episodic memory thus requires activity in cortical areas with receptive fields that can discriminate between different types of groups and different types of groups of groups of sensory objects. These types of receptive fields are polymodal and therefore located in the frontal cortex. Episodic memory will also require activity in the higher visual areas to provide receptive fields corresponding with individual objects and in the hippocampal system to manage the activation of the appropriate receptive fields in all the other cortical areas. The hippocampal activity would be expected to be targetted at specific groups of cortical columns and therefore smaller. Because episodic memory retrieval requires a specific sequence of indirect activation behaviours, cerebellar activity is therefore to be expected. Functional neuroimaging of the brain during episodic memory recall is consistent with this picture, with strong activity in the prefrontal cortex (Fletcher et al. 1997) and in visual and visual association areas (Addis et al. 2007); somewhat weaker activity in the hippocampal system (Fletcher et al. 1997); and strong cerebellar activity (Fliessbach et al. 2007).

6.6 Priming Memory

The primary information mechanism supporting priming memory is indirect activation of cortical columns on the basis of recent simultaneous receptive field detection. This will be illustrated by the example of identification of pictures following subliminal presentation. In one experiment (Badgaiyan 2000), subjects are shown a number of line drawings, studying each for 3 s. A few minutes after completing this study phase, the subjects are shown a series of pictures, each for 16 ms, and asked to identify them. The subjects had an 82% success rate for studied pictures, but only a 5% success rate for identifying new unstudied pictures.
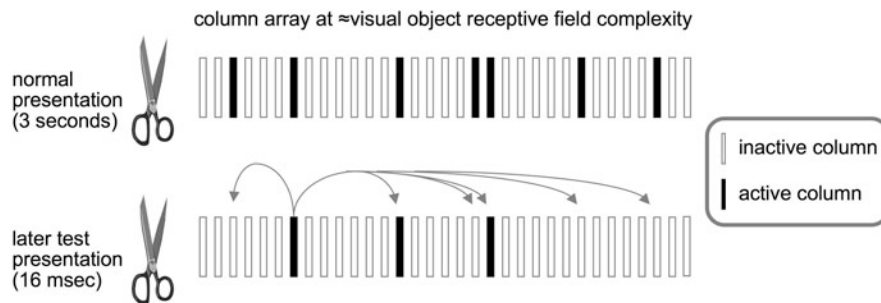

Figure 27 Column activations during study and test phases of a priming experiment. The study phase presentation results in a substantial set of column activations. A later test presentation is very brief, and only a small subset of the columns activated at study is activated. However, these columns have significant recommendation strengths in favour of activation of the other columns in the full set, on the basis of recent simultaneous activity.

As illustrated conceptually in figure 27, the study presentation results in activation of a normal sized set of columns. The very brief test presentation results on activation of a small subset of the original set that in general does not have enough total recommendation strength in favour of naming the picture. However, if the picture has

been recently studied, the columns in the subset all have strong recommendation strengths in favour of activation of the other columns in the full set on the basis of recent simultaneous activity. The resultant larger set of columns often has enough total recommendation strength to identify the picture.

Indirect activations based on recent simultaneous activity decay fairly rapidly unless repeated or used in a situation that result in reward feedback. Priming memory is the situation in which neither of these factors is present. Semantic memory can start from the same initial weights, but repetition and reward feedback strengthens and stabilizes the weights long term.

6.7 Procedural Memory
The primary information mechanism supporting procedural memory is adjustment of the behavioural recommendation weights of receptive field detections on the basis of reward feedback. In learning a new skill, if previously existing receptive fields have enough discrimination to distinguish between situations in which different skilled behaviours are appropriate, the skill could be learned by adjustment to recommendation weights in the basal ganglia without changes to cortical receptive fields. If the sensory environment in which the skill is practiced is novel, some receptive field changes will be required. The model thus accounts for the observations that patients who have lost the ability to create new declarative memories can nevertheless acquire simple mechanical skills (Corkin 1968).

The symptoms of a number of disorders associated with the basal ganglia, such as Parkinson's and Huntingdon's syndromes, can be understood in terms of the information model for the basal ganglia shown in figure 16.

The symptoms of Parkinson's disease include difficulty with voluntary movement and with initiation of movement and in general slowness of movement. The observed physical deficit (Jankovic 2007) is degeneration of dopaminergic neurons in the SNc. In the information model for the basal ganglia, this will mean that the loop that ensures that one and only one behaviour is selected in response to each sensory input state is unbalanced. The lack of dopamine will result in increased activity in the indirect pathway (figure 16), leading to no behaviour being selected.

The major symptom of Huntingdon's disease is the intrusion of irregular, unpredictable, purposeless, rapid movements that flow randomly from one body part to another (Berardelli et al. 1999). The observed physical deficit is loss of striatal cells that project into the indirect pathway (Starr et al. 2008). In the information model for the basal ganglia, this can be understood as reducing the inhibition of all behaviours. Any individual receptive field detection by the cortex will recommend a wide range of behaviours. The weakening of the selection management leads to the selection of multiple behaviours. The motor system can only implement one or a consistent set of behaviours, but the ultimate selection will be fairly random.

7 Mapping Between Different Levels of Description
As an illustration of how the recommendation architecture model makes it possible to describe the same phenomenon consistently on different levels of detail from psychology to physiology, consider again an example of episodic memory.

At the highest level, the causal description begins with a subject with a range of sensory inputs including the verbal input "What do you remember about your first day at your current job?" In the situation of the subject, the verbal input causes the subject to give priority to answering the enquiry over other possible demands. This priority causes attention to be paid to the verbal input. The verbal input causes a sequence of internal brain activities. The end point of these brain activities causes a verbal response describing the suggested event.

At a more detailed level, a range of sensory input is reaching the subject, including auditory input of the spoken words. A range of receptive fields are detected by the cortex within the words and other sensory inputs, each with a range of behavioural recommendation weights into the basal ganglia. There is substantial total recommendation weight in favour of paying attention to the auditory inputs, and as a result, the auditory information is released by the thalamus for more detailed receptive field detections by the cortex. More detailed cortical receptive fields detected within the word "remember" and receptive fields indirectly activated as a result have substantial recommendationweight in favour of an episodic memory behaviour. This weight is the current largest, and therefore a sequence of behaviours supporting episodic memory is activated in the cerebellum. This sequence drives indirect activation behaviours in the cortex, using information recorded in the hippocampal system, and favours acceptance of speech recommendationweights of the final population.

At the next level of detail, the description of the complete end to end psychological process would be very lengthy. As in the physical sciences, understanding is based on establishing more detailed descriptions of key small segments of the overall high level process.

One segment could be a description of the handling of the sequence of words making up the question. Receptive fields detected within the auditory input recommend a standard speech processing behavioural sequence recorded in the cerebellum. Such a behavioural sequence generates bias in favour of the weights of currently detected receptive fields in favour of the next behaviour in the sequence. The sequence of behaviours would be indirect activation of visual and associative receptive fields by the auditory receptive fields detected within a word, prolonging the indirectly activated population with a specific phase of frequency modulation, generating an indirectly activated population in response to a second word and prolonging it with a different phase of frequency modulation, and once several words have corresponding indirectly activated populations, bringing their outputs into the same modulation phase and releasing them to a cortical level detecting a higher level of receptive field complexity.

A second segment would be to take the population indirectly activated at the higher level of receptive field complexity by the combination of words and drive a series of further indirect activations on the basis of simultaneous past receptive field expansion, using information recorded in the hippocampal system. Again, this segment would be a sequence of behaviours instantiated in the cerebellum. The segment would include releasing the outputs of the population to the hippocampal system, where columns would detect receptive fields if a significant proportion of the columns making up their receptive fields were active. The active hippocampal columns would then encourage activation of all the cortical columns defining their receptive field. Some cortical columns could appear in multiple activated hippocampal receptive fields, and such cortical columns

would be most likely to be activated. All the releases of column outputswould bemanaged by the basal ganglia and thalamus on the basis of active column recommendation weights, modulated by the bias placed on weights in favour of certain behavioural types by the behaviour sequence instantiated in the cerebellum.

At an even deeper level, there could be descriptions of the processes for column receptive field detection, for behaviour selection through the basal ganglia and thalamus, for management of behaviour sequence through the cerebellum and for hippocampal system management of receptive field expansions.

At a yet deeper level, descriptions would be in terms of neuron receptive field definition using LTP type mechanisms, neuron receptive field detection using leaky integration and frequency modulation.

An even deeper level could describe synapses, neurotransmitters and ion channels supporting leaky integration at a higher level, and the changes to ion channels supporting LTP mechanisms.

The model thus has the capability to map consistently all the way from psychology to the chemical processes underlying physiology. This consistent mapping means that it is possible to have confidence in the intermediate level descriptions (e.g. in terms of cortical columns) which are most critical for understanding the psychological phenomena.

8 More Complex Cognitive Processes

The focus of this chapter has been on memory and learning phenomena. However, in this section an outline will be provided of how the mechanisms discussed for memory and learning relate to major cognitive management systems such as attention and emotion, and how the mechanisms supportmore complex cognitive phenomena including imagination and creativity.

8.1 Attention

The attention function on a psychological level selects a subset of the currently available sensory information and favours that subset in the determination of behaviour. Unmodulated sensory input from the whole visual field enters the cortex via the thalamus. Because the inputs are unmodulated, they do not penetrate deeply into the visual cortex. Receptive field detection is mainly in V1, where receptive fields correspond with boundary elements in different boundaries within the visual field. Each such receptive field recommends focussing attention on the retinal area in a band perpendicular to the boundary element and on both of its sides as illustrated in figure 28.

For a closed boundary, the recommendation strengths within the boundary from all the surrounding boundary elements will reinforce each other, and there will be a strong total recommendation strength in favour of the area within the boundary, in other words, an object in the visual field.

The strongest such recommendation, totalled across all currently detected receptive fields, is accepted. This acceptance is implemented by placing a modulation on the sensory inputs within the corresponding closed boundary. The effect of the modulation is that receptive fields are detected deep into the visual cortex and beyond. These receptive field detections recommend behaviours appropriate to the visual object within the selected closed boundary.
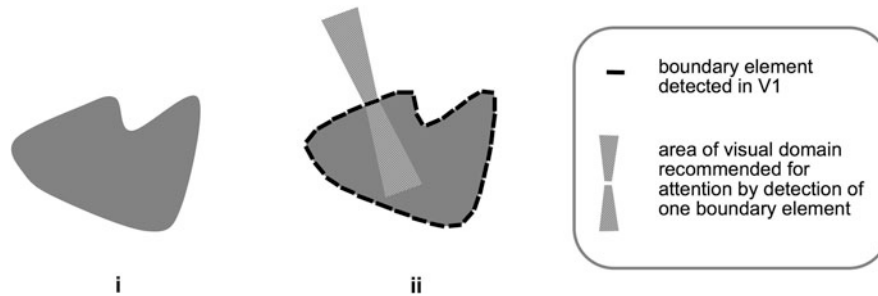
Figure 28 Management of attention by boundary element detection in V1. In (i), a visual field is illustrated with one object. In (ii), some of the boundary element receptive fields detected in visual area V1 within the object illustrated in (i) are illustrated, along with the retinal area recommended for attention by one receptive field detection.Within a closed boundary, attention recommendations will reinforce each other, resulting in a strong overall recommendation in favour of attention on the area within the closed boundary.

Receptive field detections within multiple objects can be retained active without interference with each other by modulating the detections at different phases of modulation as illustrated in figure 13. This support of multiple sets of receptive field detections supports working memory as discussed earlier.

Analogous mechanisms support attention in auditory and propioceptic sensory processing.

8.2 Emotion and Reward

Different cortical areas detect receptive fields within different ranges of complexity. Because of the need to economize on resources, receptive field detections in any area recommend many different behaviours. However, as discussed earlier, there may be behavioural advantages in having areas that specialize in receptive fields effective in discriminating between circumstances in which different behaviours of a particular general type (e.g. food seeking, aggressive, avoidance, etc.) are appropriate. Sometimes the behavioural advantage of such areas may outweigh the extra resource cost. Such areas have the additional advantage that the relative probability of their behaviour type can be modulated.

Emotions correspond with different general types of behaviour, and the presence of the emotion encourages its corresponding type. Anger encourages aggressive behaviour, fear encourages avoidance (i.e. fearful) behaviour, disgust encourages rejection behaviour, surprise encourages recording information and delaying behaviour, sadness encourages avoidance of behaviour because of radical change in circumstances, happiness encourages repetition of recent behaviours and hunger encourages food seeking behaviour, etc.

Some cortical receptive field detections recommend such general types of behaviour. These recommendations are instantiated by connections to the amygdala and/or hypothalamus, with different weights. If the total recommendation weight in favour of a behaviour type is high enough, the amygdala and/or hypothalamus will be activated to release a neurotransmitter to targets in the cortex within areas that generate more specific behavioural recommendations of the general type. Such a neurotransmitter release corresponds with an emotional state and will have the effect of modulating the thresholds

of pyramidal neurons in the targetted areas. This modulation will result in much more activity recommending specific behaviours of the general type.

Emotion signals generated by the amygdala and hypothalamus also target the striatum, where they influence the relative probability of a behaviour of their corresponding type being accepted. Thus, for example, the basolateral complex of the amygdala, which is associated with fear, projects strongly to both the striatum and the prefrontal cortex (Sah et al. 2003).

Note that if a set of columns recommending an emotion are active at the same time as a set of cortical columns elsewhere in the cortex, the two sets are active at the same time and can therefore acquire recommendation strengths in favour of activating each other. At a psychological level, this means that, for example, the memory of an event may generate an emotion.

Note also that the generation of a reward that modulates recent recommendation strengths is itself a behaviour that is in general recommended by cortical receptive field detections. If a set of columns are active at the same time as some reward recommending columns, the set may acquire recommendation strengths in favour of indirect activation of the reward columns. The circumstances in which a particular type of behaviour is appropriate will be a very complex combination of sensory inputs able to discriminate between subtle social differences. Reward discrimination in many cases also requires subtle discrimination between different social circumstances. The receptive field complexity for the two types of discrimination may be comparable, resulting in the same cortical area being associated with both types of behavioural recommendation. The orbitofrontal cortex does appear to be involved in both emotion and rewards (O'Doherty et al. 2001).

8.3 Sleep

A major role of sleep in the recommendation architecturemodel is the configuration of neural resources that are as appropriate as possible for learning in subsequent waking periods (Coward 1990, 2001). Recording of information in the brain occurs by creation of new synapses and by changes to synaptic weights. If new synaptic connections are needed, it would not be practical to create those synapses at the instant they were required. It is therefore necessary to create "provisional" synapses in advance, which can be utilized if appropriate during subsequent experience. Such provisional synapses could be between randomly selected pre- and postsynaptic neurons, but this would result in considerable resource expenditure on the creation of synapses which are not useful and which in some circumstances could reduce behavioural effectiveness (Coward 2001).

One way to improve the probable effectiveness of provisional synapses is to utilize past experience. If two neurons have never in the past been active at the same time, the probability that a connection between them will be useful is low. Conversely, if two unconnected neurons have often been active at the same time in the past, the probability is higher. If the simultaneous activity is recent, the probability is even higher. The strategy is therefore to bias the creation of provisional synapses in favour of connections between neurons that have tended to be active at similar times in the past, with the recent past having the strongest weight. Provisional connectivity will also be required to support indirect activation on the basis of past temporally correlated activity, such as activity of one neuron at the same time as or shortly after another neuron.

Connectivity is required between cortical neurons and between cortical and hippocampal neurons to support declarative memory. Connectivity is also required between cortical neurons and basal ganglia and thalamic neurons, and within the basal ganglia and thalamus, to support procedural memory. Different stages of sleep may support different aspects of provisional connectivity creation, including extension of axon segments and building of synapses and connectivity in different brain structures. Partial reruns of past experience will be required to identify the best candidates for provisional connectivity.

Consistent with the model, reruns of past patterns of neuron activity have been observed in both REM sleep (Pavlides and Winson 1989) and in slow wave sleep (SWS) (Lee andWilson 2002). There is some evidence that reruns of recent experience occur in SWS, with reruns of more remote experience in REM sleep (Hoffman and McNaughton 2002).

This model is radically different from consolidation models (Squire and Alvarez 1995) in which memory information is initially recorded in the hippocampus and gradually over time transferred to the cortex. Such information transfers would be very complex, and modelling efforts have generally been limited to just the hippocampus (e.g. Redish and Touretzky, 1998). The resource configuration model does not require these complex information transfers, memories are fully defined in cortical information terms at the end of the few hundred millisecond period of the remembered experience. There may be chemical processes required to consolidate the memory (Tronson and Taylor 2007), but the information content of the memory is not significantly changed or relocated. The hippocampus is only required in the future to identify the sets of cortical information associated with individual episodic memories. Furthermore, the rerun process is not indispensible for declarative memories, since random connectivity can support such memories with reduced resource effectiveness. The resource configuration model is therefore more consistent with the observation that substantial or complete suppression of REM sleep by various antidepressant drugs or by bilateral damage to the pons has no apparent effect on declarative memory (Vertes and Eastman 2000).

8.4 Mental Image Manipulation

Consider how a brain might be able to imagine the appearance of a currently viewed object as it would appear if viewed from a different angle. As an example, suppose that a subject is presented with an image of an object viewed from an unfamiliar angle and is asked to identify the object. The requirement is to be able to mentally rotate the current visual image to generate a mental image that is more recognizable.

In the subject's past experience, many different objects have been examined from many different angles. During these examinations, the cortical columns activated in response to viewing from one angle have been active shortly before or after the columns activated in response to viewing from a different angle. As a result, columns will have acquired recommendation strengths in favour of activating other columns on the basis of frequent past activity just before and just after the other columns.

The behaviour of "mentally rotating" is therefore one of imposing a bias on the behavioural selection process in favour of column recommendations of indirect activation on the basis of frequent past activity just before and just after the other columns. Any one column will have been activated in the past in response to many different objects. There

will be a population of columns currently active in response to the visual image. The effect of this recommendation type will be activation of a population of columns, each of which has often been active in the past just before or after a number of these currently active columns. This indirectly activated (or secondary) populationwill be maintained active at a different phase of frequency modulation from the directly activated columns, to avoid meaningless mixing of information.

The secondary population will therefore correspond with an estimate of what the object would look like from a different angle based on an average of many past experiences of viewing rotations of different objects. The recommendation strengths of this population in favour of naming an object will be accepted. This process is of course not guaranteed to work, since it depends upon the existence of reasonably frequent, relevant past experiences.

A key point is that the brain does not "know about" rotations, but can implement processes based on indirect column activations. Certain processes generate results that are rewarded. For example, when initially a child is asked to imagine an object as it would be if rotated, various indirect activation processesmay be tried randomly. Once the appropriate indirect activation process is tried, the result will be rewarded, and the process will tend to be used again in the future if the word "rotated" is heard.

8.5 Self-Awareness

Consider a child who is learning his name, and suppose that name is Michael. The name is used on many occasions when his attention is being directed towards himself. When his attention is directed towards himself, cortical columns detect their receptive fields within visual information derived from looking at himself and within proprioceptic information derived from his own body ("that's Michael's foot"). Other columns may be active recommending various emotions. In addition, the word "Michael" activates auditory column receptive fields.

Over a period of time, there will be some columns (visual, proprioceptic, polymodal, emotional, etc.) that are often active at the same time as a group of auditory columns often active when the word "Michael" is spoken. Hence the auditory columns will acquire recommendation strengths in favour of indirect activation of the visual, proprioceptic, polymodal and emotional columns. The result is that the auditory columns activate an internal state that is a kind of average of Michael's experiences at the times when his name has been used.

Michael has often heard the word "boy" when his visual attention has been directed towards various boys, and frequent simultaneous column activity means that the columns activated in response to hearing the word will tend to activate a set of visual columns that are an average of the boys Michael was looking at when he heard the word.

Furthermore, Michael sometimes hears the words "Michael is a boy". At this point, the visual columns indirectly activated by hearing "boy" are active at the same time as the auditory columnswithin the word "Michael". Hence those auditory columns will acquire recommendation strengths in favour of activating "boy" visual columns. Finally, Michael sometimes looks in a mirror and is told "that's Michael", and again the auditory receptive fields detected within the word are active at the same time as the visual receptive fields detected within the reflection of Michael.

The result is that hearing the word "Michael" can result in activation of a population of columns that are a kind of average of Michael's emotions, visual inputs and body awareness when his attention was on himself viewed from the "inside", plus a population of columns that are a kind of average of Michael's sensory experiences of other boys and himself viewed from the "outside". These two self images will generally be activated at different phases of frequency modulation from each other and from current sensory inputs, etc. so that they will not be confused. The "inside" and "outside" images correspond with the "me" and the "I" perspectives (Jaynes 1976), the associations with the "me" and "I" words being again created by frequent simultaneous activity of columns.

8.6 Imagination

Suppose that someone is asked "imagine that you are at a party with Bill Clinton". The word "you" drives activation of a self column population as described in the previous section. This activation could be at many different levels of column receptive field complexity. The words "Bill Clinton" drive activation of a column population on the basis of frequent past simultaneous activity with the auditory columns, corresponding with mental images that are an average of past visual observations of Bill Clinton on television, columns activated in response to stories about Bill Clinton, etc. The word "party" will similarly activate a population of columns including, for example, columns active during past parties actually attended.

There will be strong active column populations in the visual objects and groups of objects cortical areas, because self is familiar, Bill Clinton and some actions of Bill Clinton are familiar, and parties are familiar. However, the populations in the groups of objects and especially in the groups of groups, etc. areas will have some weaknesses, because no party including self and Bill Clinton has been attended. These weaknesses will drive receptive field expansions, which because they are adding conditions containing information derived from self, parties and Bill Clinton will be the type of expansions that would occur if a party with Bill Clinton present was attended. However, there will be little related activity in the cortical areas very close to sensory inputs, so the imagining will not be a visual hallucination. The result will therefore be an experience as if recalling a party with Bill Clinton.

Because there has been receptive field expansion, it will be possible to recall imagining the event. However, in such a recall the level of activity close to sensory inputs will be even less than in recall of a real experience, so in general it will not be confused with a real memory.

Initially, the imagining will be at a relatively high level of receptive field complexity, and therefore experienced as a relatively abstract experience. However, the active receptive fields will have recommendation strengths to activate other receptive fields on the basis of frequent past simultaneous activity. Some simpler receptive fields closer to sensory input will have been active in the past at the same time as the higher level receptive fields and will therefore be activated. If the degree of activity at the simpler level is small, there could be receptive field expansion at this simpler level. Hence a less abstract experience can be generated. It is possible that such a process could lead to an imagined experience that in retrospect was indistinguishable from a real experience, in other words a false memory as observed by Loftus and Pickrell (1995).

This model is consistent with the observations that the brain regions activated in response to recalling a real event are very similar to the regions activated in response to imagining an event that has never happened (Addis et al. 2007).

The capability to imagine events well outside of actual past experience at anything other than a very abstract level is one aspect of what is called creativity. This capability will depend on whether some required receptive field expansions are possible.

The column activations that support detailed imagination are indirect activations of simpler receptive fields on the basis of past simultaneous activity with more complex receptive fields that are already active.

For a simpler receptive field to be activated, it must have been active in the past at the same time as a significant number of the more complex receptive fields. For an imagined experience that is radically different from any past experience, there will be relatively few simpler receptive fields that have the required level of temporally correlated past activity. Hence the total activity in the cortical area within which the columns with the simpler receptive fields are located will be lower than the minimum required level. The inputs driving activation in this area are not regular sensory but indirect activation inputs. Receptive field expansions could occur to increase the degree of area activity, but these expansionswould be new groupings of currently active columns at the more complex level. In information terms, this introduces a kind of semantic memory based not on actual experience but on imagined experience.

Receptive field expansions can only occur if at the pyramidal neuron level there are adequate provisional conditions. In this case, such provisional conditions would be combinations of higher complexity columns supporting indirect activation. Such provisional conditions would need to be created during sleep processing.

There are behavioural benefits and costs to the creation of such provisional conditions, which are essentially opposite sides of the same coin. The benefit is the ability to imagine circumstances well outside past experience, and therefore, for example, the ability to plan for such unusual circumstances. The cost is the introduction of indirect activation capabilities that are not supported by actual experience and may be misleading for future planning. The degree to which a brain supports such provisional condition configuration in the hippocampal system, the anterior temporal cortex and the monomodal sensory and polymodal cortices is therefore a compromise which may be different for different brains. This compromise defines the degree of creativity which the brain will support.

8.7 Planning

To illustrate planning, suppose that someone is asked "What will you do after the concert?" The words "you" and "concert" activate a primary population of auditory cortical columns, which in turn indirectly activate a secondary population of visual and polymodal columns on many levels of receptive field complexity. Hearing the words "what", "do" and "after" encourages favouring a behaviour sequence recorded in the cerebellum. This sequence biases behaviour acceptance in the basal ganglia and thalamus in favour of a sequence of different indirect activation behaviours.

The first step in the sequence biases acceptance of recommendation strengths in the secondary population of columns in favour of activation of other columns often active in the past after the columns active in the secondary population were active. The effect will

be to activate a tertiary population made up of fragments of populations that were active during different past "after concert" experiences.

The next step in the sequence biases acceptance of recommendation strengths in the tertiary population in favour of activation of other columns that expanded their receptive fields in the past at the same time as the columns active in the tertiary population. This step also preserves the activity of the tertiary population, at a different modulation phase from the developing quaternary population. The effect is to establish a quaternary population approximating to the population active during the event corresponding with the fragment with the strongest total recommendation strengths in the tertiary population. In other words, this quaternary population corresponds with the past "self after concert" experience with the largest active column representation in the tertiary population.

The next step in the sequence biases acceptance of recommendation strengths in the quaternary population in favour of activation of other columns that were often active in the past after the columns in the quaternary population. The effect is to establish a quinary population approximating to the population active a little later than the population active during the "self after concert" past experience. If this population contains a lot of columns active in the past at the same time as columns recommending a positive reward, the quaternary population becomes the "plan" for after concert.

If there is little positive reward recommendation in the quinary population, the brain can go back to the tertiary population (made up of fragments of different "after concert" experiences) and generate another population by activation of other columns that expanded their receptive fields in the past at the same time as the columns active in the tertiary population. However, in this case, there is a bias placed against recently active columns, resulting in the activation of a population corresponding with the past "self after concert" experience with the second largest active column representation in the tertiary population. A further population is indirectly activated to test reward level after the new quaternary population.

This tertiary to quaternary to quinary process can be repeated until a high reward level is found in a quinary population, at which point the final quaternary population becomes the "plan". The final step is to repeat the evolution from the "self after concert" population to the final quaternary population.

Note that for simplicity, more complex sequences involving imaginary scenarios developed in a similar way to that described in the previous section could also be developed and evaluated for overall final positive reward recommendation strengths. Also, in reality there will be intermediate populations approximating to different stages in past experiences. Such intermediate populations can be added to the above description without difficulty.

At the end of the concert, a population with a fair amount of overlap with the tertiary population generated earlier will develop. Because the populations corresponding with the plan will have recently been active shortly after that population, acceptance of recommendation strengths in favour of activation on the basis of recent activity shortly after the currently active columns will result in development of the column populations corresponding with the plan.

All active column populationswill also have recommendation strengths in favour of motor behaviours to get to the next stage in the experience. However, during the concert, these recommendation strengths will not be the strongest in total. In the new sensory

circumstances corresponding with the end of the concert, these motor behaviours will have the predominant total recommendation strengths.

8.8 Stream of Consciousness

Perhaps the classical definition of human consciousness is the experience of a stream of mental images relatively unrelated to current sensory inputs. According to James, the experience is of a stage with relatively consistent mental images, separated by periods of vague evolution (James 1892).

In the recommendation architecture, the starting point for this process is some population of columns directly activated by the presence of their receptive fields within current sensory inputs. Each activated column has a set of recommendation strengths in favour of activating other columns on the basis of past temporally correlated activity which could be recent activity at the same time, before or after, frequent past activity at the same time, before or after, or past receptive field expansion at the same time, before or after.

If such indirect expansions are encouraged, in the absence of any strong total recommendation in favour of an externally directed behaviour, a vague secondary population will develop. Evolution could then continue through tertiary, quaternary, etc. populations, but because of the huge range of indirect activation recommendation strengths the cognitive content of the population could become extremely vague.

However, there is a way to focus a population at any point in time to make it potentially more cognitively useful (e.g., for developing a plan as discussed earlier). Any population will also have recommendation strengths in favour of speech. If the strongest such recommendation is accepted, but only to activate auditory columns often active in the past before the columns in the population were active, the experience will be of pseudohearing of words that are the closest representation of the population. If now the population is replaced with another population generated on the basis of frequent past activity at the same time as the auditory columns, the effect will be a population with a more sharply focussed cognitive meaning.

Evolution of the new population could then occur, eventually followed by another focussing stage. This model thus accounts for the stream of consciousness experience as described by James. For more detail see Coward and Gedeon (2009).

9 Electronic Implementations

The recommendation architecture memory model is dependent upon a number of information mechanisms. One of the most critical is unsupervised organization of experience into column condition groups, in such a way that a set of columns can discriminate between different circumstances if the difference is behaviourally significant. A second mechanism is association of different groups of columns with different appropriate behaviours using reward feedback, in such a way that interference between new and prior learning is minimized. A third mechanism is support for different independent populations of active columns within the same resources, using different phases of frequency modulation. A fourth is indirect activation of columns on the basis of past temporally correlated activity. A fifth is management of the configuration of provisional connectivity for future receptive field expansions using past experience.

There have been various electronic implementations of these information mechanisms that confirm their capabilities. These implementations have in general used three layer columns as illustrated in figure 2, pyramidal neuron receptive fields that can expand slightly but not contract or change qualitatively, and have used software (Smalltalk and CCC) models for physiological structures. Early implementations (e.g. Gedeon et al. 1999; Coward 2001; Ratnayake et al. 2003; Coward et al. 2004) have used a relatively simple pyramidal neuron model in which all synapses have the same weight and inputs from one sensory state arrive synchronously. Later neuron models (e.g. Coward 2004, 2009b) are dynamic, with staged leaky integration across their dendritic trees, and use an LTP algorithm for learning.

These electronic simulations have demonstrated that experience can be organized into column modules, where each column detects a gradually expanding similarity space that is relatively orthogonal to the spaces detected by other columns, in such a way that the column ensemble can discriminate between circumstances with behaviourally different implications (Gedeon et al. 1999; Coward 2001, 2009b; Ratnayake et al. 2003). The ability to associate partially ambiguous columns with behaviours using reward feedback and the capability of imitation to improve the efficiency of rewardbased learning (Coward 2005) have been demonstrated, including the management of behavioural selection by competition between components corresponding with the different behaviours (Coward et al. 2004). It has also been demonstrated that the gradual expansion of column portfolios means that the architecture does not experience catastrophic interference (Coward et al. 2004). Behaviours have included appropriate responses to objects and groups of objects. Simulations have also demonstrated the effectiveness of indirect activation mechanisms in supporting activation of pseudovisual images in response to verbal inputs and supporting activation of pseudovisual images of objects often present in the past at the same time as currently perceived objects (Coward 2001). The capability of the frequency modulation mechanism to implement attention functions at the physiological level has also been demonstrated (Coward 2004). The use of a sleep-like process for configuration of provisional conditions has also been implemented, with the expected improvement to the behavioural effectiveness of recorded conditions (Coward 2001).

10 Conclusions
The recommendation architecturebased cognitive model establishes the consistent causal descriptions of memory and learning phenomena on different levels of detail that are the essential core of any scientific understanding. Causal descriptions of memory and learning at the psychological level can be precisely mapped into causal descriptions of the same phenomena at the level of anatomical structures. The anatomical descriptions can be precisely mapped into causal descriptions at the level of neuron algorithms, and the neuron level descriptions can be mapped into known physiology. Alternative theories do not demonstrate the same degree of consistent multilevel modelling.

A critical intermediate level of description is at the level of cortical columns, because the most complex types of information processing occur in the cortex. However, units of cortical information processing cannot be mapped simply into units of cognitive processing. The units of information processing in the cortex are detections of similarity circumstances (or receptive fields) and expansion of receptive fields when necessary to

reach a minimum required level of detection. One such unit of information processing is shared by many different cognitive processes. All cognitive processing can be described in terms of direct detection of receptive fields, indirect activation of receptive fields on the basis of past temporally correlated activity with currently active receptive fields and interpretation of receptive field detections as behavioural recommendations, including indirect activation recommendations.

The brain has no a priori definitions of cognitive processes like episodic memory, imagining, planning, etc. All the brain has is a set of information processes (or internal behaviours) for activating and evolving populations of receptive field detections. At each stage the currently active population is interpreted into a predominant behaviour recommendation which is then implemented. Learning uses reward feedback both to discover sequences of information processes that are behaviourally valuable and to associate different circumstances with invocation of different sequences. Different sequences supporting different cognitive processes such as episodic memory, semantic memory and imagination or planning often use overlapping cortical resources. Electronic modelling of the recommendation architecture demonstrates that the information models that form the foundation of the description hierarchy support observed memory and learning phenomena.

References

Addis DA, Wong AT, Schacter DL (2007) Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. Neuropsychologia 45:1363–1377.

Alexander GE, DeLong MR, Strick PL (1986) Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annual Reviews of Neuroscience 9:357–381.

Anderson AK, Phelps EA (1997) Emotional memory: what does the amygdala do? Current Biology 7:R311–R314.

Anderson JR (1996) ACT: A simple theory of complex cognition. American Psychologist 51:355–365.

Axmacher N, Morman F, Fernandez G, Elger C, Fell J (2006) Memory formation by neuronal synchronization. Brain Research Reviews 52:170–182.

Baddeley AD (1986) Working Memory. Oxford, Oxford University Press.

Baddeley AD, Vallar G, Wilson BA (1987) Sentence comprehension and phonological memory: some neuropsychological evidence. In: Attention and Performance XII: The psychology of reading. London: Erlbaum.

Baddeley AD (2000) Short-term and working memory. In E. Tulving and F. I. M. Craik (Eds.). The Oxford Handbook of Memory. Oxford: Oxford University Press.

Badgaiyan RD (2000) Neuroanatomical organization of perceptual memory: An fMRI study of picture priming. Human Brain Mapping 10:197–203.

Bainbridge L (1977) Verbal reports as evidence of the process operator's knowledge International. Journal of Man-Machine Studies 11:411–436.

Bar M, Biederman I (1998) Subliminal visual priming. Psychological Science 9:464–469.

Berardelli A, Noth J, Thompson PD et al. (1999) Pathophysiology of chorea and bradykinesia in Huntington's disease. Movement Disorders 14:398–403.

Berry DC (1987) The problem of implicit knowledge. Expert Systems 4:144–151.

Bi G-Q, Poo M-M (1998) Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. Journal of Neuroscience 18:10464–10472.

Bracke-Tolkmitt R, Linden A, Canavan AGM et al. (1989) The cerebellum contributes to mental skills. Behavioral Neuroscience 103:442–446.

Brodmann K (1908) Beitraege zur histologischen Lokalisation der Grosshirnrinde. VI. Mitteilung: Die Cortexgliederung desMenschen. Journal of Psychology and Neurology (Lzp) 10:231–246T.

Cabeza R, Dolcos F, Graham R et al. (2002) Similarities and differences in the neural correlates of episodic memory retrieval and working memory. Neuroimage 16:317–30.

Cohen JJ, Squire LR (1981) Preserved learning and retention of pattern analysing skill in amnesia: Dissociation of knowing how and knowing that. Science 210:207–210.

Collins A, Loftus E (1975) A spreading-activation theory of semantic processing. Psychological Review 82:407–428.

Corkin S (2002) What's new with the amnesic patient H.M.? Nature Reviews Neuroscience 3:153-160

Corkin S (1968) Acquisition of motor skill after bilateral medial temporal-lobe excision. Neuropsychologia 6:225–264.

Cowan N (2000) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. The Behavioural and Brain Sciences 24:87–185.

Coward LA (1990) Pattern Thinking. New York, Praeger.

Coward LA (2001) The Recommendation Architecture: lessons from the design of large scale electronic systems for cognitive science. Journal of Cognitive Systems Research 2:111–156.

Coward LA (2004) Simulation of a proposed binding model. In: Brain Inspired Cognitive Systems 2004, Smith LS, Hussain A, Aleksander I (eds). Stirling, University of Stirling.

Coward LA, Sun R (2004) Some criteria for an effective scientific theory of consciousness and examples of preliminary attempts at such a theory. Consciousness and Cognition 13:268–301.

Coward LA, Gedeon TD, Ratanayake U (2004) Managing interference between prior and later learning. ICONIP 2004, Calcutta. Lecture Notes in Computer Science 3316:458–464.

Coward LA (2005) A System Architecture Approach to the Brain: from Neurons to Consciousness. New York, Nova Science Publishers.

Coward LA, Sun R (2007) Hierarchical approaches to understanding consciousness. Neural Networks 20:947–954.

Coward LA (2009a) The Hippocampal system as the cortical resource manager: A model connecting psychology, anatomy and physiology. In: Brain Inspired Cognitive Systems. Cutsuridis V, Hussain A, Barros AK, Aleksander I, Smith L, Chrisley R. (eds). Berlin, Springer.

Coward LA (2009b) The Hippocampal system as the manager of neocortical declarative memory resources. In: Connectionist Models of Behaviour and Cognition II, Mayor J, Ruh N, Plunkett K (eds). London, World Scientific.

Coward LA, Gedeon TO (2009) Implications of resource limitations for a conscious machine. Neurocomputing 72:767–788.

Crovitz HF, Schiffman H (1974) Frequency of episodic memories as a function of their age. Neuropsychologia 21:213–234.

Crutcher MD, DeLong MR (1984) Single cell studies of the primate putamen II. Relations to direction of movement and pattern of muscular activity. Experimental Brain Research 53:244–258.

Cutsuridis V, Cobb S, Graham BP (2010) Encoding and retrieval in the hippocampal CA1 microcircuit model. Hippocampus 20(3):423–446.

Cutsuridis V, Wenneckers T (2009) Hippocampus, microcircuits and associative memory. Neural Networks 22(8):1120–1128.

Cutsuridis V, Cobb S, Graham BP (2008) Encoding and retrieval in a CA1 microcircuit model of the Hippocampus. In: Lecture Notes in Computer Science 5164, Kurkova V, et al. (eds). Berlin, Springer, 238–247.

Devinsky O, Morrell MJ, Vogt BA (1995) Contributions of anterior cingulate cortex to behaviour. Brain 118:279–306.

Devlin JT, Russell RP, DavisMH et al. (2002) Is there an anatomical basis for category-specificity? Semantic memory studies in PET and fMRI. Neuropsychologia 40:54–75.

Engel AK, Singer W. (2001) Temporal binding and the neural correlates of sensory awareness. Trends in Cognitive Sciences 5:16-25.

Fan Y, Zou B, Ruan Y et al. (2005) In vivo demonstration of a late polarizing postsynaptic potential in CA1 pyramidal neurons. Journal of Neurophysiology 93:1326–1335.

Fliessbach K, Trautner P, Quesada CM et al. (2007) Cerebellar contributions to episodic memory encoding as revealed by fMRI. NeuroImage 35:1330–1337.

Fletcher PC, Frith CD, Rugg MD (1997) The functional neuroanatomy of episodic memory. Trends in Neurosciences 20:213–218.

Fyhn M, Molden S, WitterMP, Moser EI, Moser M-B (2004) Spatial representation in the entorhinal cortex. Science 305:1258–1264.

Gabrieli JDE, Milberg W, Keane MM et al. (1990) Intact priming of patterns despite impaired memory. Neuropsychologia 28:417–427.

Gabrieli JDE, Corkin S, Mickel SF et al. (1993) Intact acquisition and long-term retention of mirror tracing skill in Alzheimer's disease and in global amnesia. Behavioral Neuroscience 107:899–910.

Gedeon T, Coward LA, Bailing Z (1999) Results of Simulations of a System with the Recommendation Architecture. Proceedings of the 6th International Conference on Neural Information Processing I:78–84.

Graf P, Squire LR, Mandler G (1984) The information that amnesic patients do not forget. Journal of Experimental Psychology: Learning, Memory and Cognition 10:164–178.

Guillery RW, Feig SL, Lozsadi DA (1998) Paying attention to the thalamic reticular nucleus. Trends in Neuroscience 21:28–32.

Gulledge AT, Stuart GJ (2003) Excitatory Actions of GABA in the Cortex. Neuron 37, 299–309.

Hasselmo M, Bodelon C, Wyble B (2002) A proposed function of the hippocampal theta rhythm: Separate phases of encoding and retrieval of prior learning. Neural Computing 14:793–817.

Hausser M,Mel B (2003) Dendrites: bug or feature? Current Opinion in Neurobiology 13:372–383.

Heindel WC, Salmon DP, Shults CW et al. (1989) Neuropsychological evidence for multiple implicit memory systems: A comparison of Alzheimer's, Huntington's, and Parkinson's disease patients. The Journal of Neuroscience 9:582–587.

Hoffman KL, McNaughton BL (2002) Sleep on it: cortical reorganization after the fact. Trends in Neuroscience 25(1):1–2.

James W(1892) The stream of consciousness. Psychology ChapterXI.World Publishing Company, Cleveland and New York /http://psychclassics.yorku.ca/ James/jimmy11.htm

Jankovic J (2008) Parkinson's disease: clinical features and diagnosis. Journal of Neurology, Neurosurgery and Psychiatry 79:368–376.

Jaynes J (1976) The Origin of Consciousness in the Breakdown of the Bicameral Mind. Boston, Harvard.

Kassubek J, Schmidtke K, Kimmig H et al. (2001) Changes in cortical activation during mirror reading before and after training: an fMRI study of procedural learning. Cognitive Brain Research 10:207–217.

Kensinger EA, Ullman MT, Corkin S (2001) Bilateral medial temporal lobe damage does not affect lexical or grammatical processing: Evidence from amnesic patient H.M. Hippocampus 11:347–360.

Kiehl KA, Liddle PF, Hopfinger JB (2000) Psychophysiology 37:216–223.

Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. Neuron 36:1183–1194.

Leiner HC, Leiner AL, Dow RS (1993) Cognitive and language functions of the human cerebellum. Trends in Neurosciences 16:444–447.

Leutgeb S, Leutgeb JK, Treves A, Moser M-B, Moser EI (2004) Distinct ensemble codes in Hippocampal areas CA3 and CA1. Science 305:1295–1298.

Linden DJ (2003) From molecules to memory in the Cerebellum. Science 301:1682–1685.

Lisman JE, Idiart MAP (1995) Storage of 7 ± 2 short term memories in oscillatory subcycles. Science 267:1512–1515.

Loftus EF, Pickrell JE (1995) The formation of false memories. Psychiatric Annals 25:720–725.

Lovett MC, Reder LM, Lebiere C (1999) Modeling working memory in a unified architecture: An ACT-R perspective. In: Models of Working Memory, Miyake A, Shah P (eds). Cambridge, Cambridge MA, 135–182.

Mandler G, Shebo BJ (1982) Subitizing: An analysis of its component processes. Journal of Experimental Psychology: General 111:1–22.

Martin A (2007) The representation of object concepts in the brain. Annual Review of Psychology 58:25–45.

Martin SJ, Morris RGM (2002) New life in an old idea: the synaptic plasticity and memory hypothesis revisited. Hippocampus 12:609–636.

Mathews R, Buss R, Stanley W et al. (1989) Role of implicit and explicit processes in learning from examples: a synergistic effect. Journal of Experimental Psychology: Learning, Memory and Cognition 15:1083–1100.

McBride DM, Dosher BA, Gage NM (2001) A comparison of forgetting for conscious and automatic memory processes in word fragment completion tasks. Journal of Memory and Language 45:585–615.

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

Mitchell SJ, Richardson RT, Baker FH et al. (1987) The primate globus pallidus: neuronal activity related to direction of movement. Experimental Brain Research 68:491–505.

Morasan P, Rademacher J, Schleicher A et al. (2001) Human primary auditory cortex: Cytoarchitectonic subdivisions and mapping into a spatial reference system. NeuroImage 13:684–701.

Mountcastle VH (1997) The columnar organization of the neocortex. Brain 120:701– 722.

Muller MM, Gruber T, Keil A (2000) Modulation of induced gamma band activity in the human EEG by attention and visual information processing. International Journal of Psychophysiology 38:283–299.

Nadel L, Moscovitch M (1997) Memory consolidation, retrograde amnesia and the hippocampal complex. Current Opinion in Neurobiology 7:217–227.

O'Doherty J, Kringelbach ML, Rolls ET et al. (1991) Abstract reward and punishment representations in the human orbitofrontal cortex. Nature Neuroscience 4:95–102.

O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3:317–330.

Pavlides C, Winson J (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. Journal of Neuroscience 9:2907–2918.

Petrides M, Pandya DN (1999) Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque ventrolateral prefrontal cortex and corticocortical connection patterns. European Journal of Neuroscience 11:1011–1036.

Petrides M, Pandya DN (2002) Comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns in the monkey. European Journal of Neuroscience 16:291–310.

Ratnayake U, Gedeon TD (2003) Extending The Recommendation Architecture Model for Text Mining. International Journal of Knowledge-Based Intelligent Engineering Systems 7:139-148.

Redish AD, Touretzky DS (1998) The role of the hippocampus in solving the Morris water maze. Neural Computation 10:73–111.

Rips L, Shoben J, Smith E (1973) Semantic distance and verification of semantic relations. Journal of Verbal Learning and Verbal Behaviour 12:1–20.

Ritz M (1999) Chapters 5 and 6. In: Drugs of Abuse and Addiction: Neurobehavioral Toxicology, Niesink R, Jaspers RMA, Hollinger MA et al. (eds). Boca Raton, CRC.

Robinson JA (1976) Sampling autobiographical memory. Cognitive Psychology 8:578–595.

Rogers TT, Hocking J, Noppeney U et al. (2006) Anterior temporal cortex and semantic memory: reconciling findings from nuropsychology and functional imaging. Cognitive, Affective and Behavioural Neuroscience 6:201–213.

Sagar JH, Cohen NJ, Corkin, S et al. (1985) Dissociations among processes in remote memory. Annals of the New York Academy of Science 444:533–535.

Sah P, Faber ESL, De Armentia ML et al. (2003) The amygdaloid complex: Anatomy and physiology. Physiological Review 83:803–834.

Schacter JC, Tulving E (1994) What are the memory systems of 1994? In: Memory Systems 1994, Schacter JC, Tulving E (eds). Cambridge, MA, MIT.

Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. Journal of Neurology, Neurosurgery, and Psychiatry 20:11–21.

Shallice T, Warrington EK (1970) Independent functioning of verbal memory stores: a neuropsychological study. Quarterly Journal of Experimental Psychology 22:261–273.

Somogyi P, Bolam JP, Smith AD (1981) Monosynaptic cortical input and local axon collaterals of identified striatonigral neurons. A light and electron microscopic study using the Golgiperoxidase transport-degeneration procedure. Journal of Comparative Neurology 195:567–584.

Squire LR, Alvarez P (1995) Retrograde amnesia and memory consolidation: a neurobiological perspective. Current Opinion in Neurobiology 5:169–177.

Standing L, Conexio J, Haber RN (1970) Perception and memory for pictures: single-trial learning of 2500 visual stimuli. Psychonomic Science 19:73–74.

Starr PA, Kang GA, Heath S, et al. (2008) Pallidal neuronal discharge in Huntington's disease: support for selective loss of striatal cells originating the indirect pathway. Experimental Neurology 211:227-33

Sun R, Peterson T, Merrill E (1996) Bottom-up skill learning in reactive sequential decision tasks. Proceedings of 18th Cognitive Science Society Conference. Lawrence Erlbaum Associates, Hillsdale, NJ.

Sun R, Coward LA, Zenzen MJ (2005) On levels of cognitive modeling. Philosophical Psychology 18:613–637.

Tanaka K. (2003) Columns for complex visual object features in the inferotemporal cortex: clustering of cells with similar but slightly different stimulus selectivities. Cerebral Cortex 13:90–99.

Tepper JM, Bolam JP (2004) Functional diversity and specificity of neostriatal interneurons. Current Opinion in Neurobiology 14:685–692.

Torriero S, Oliveri M, Koch G et al. (2007) Cortical networks of procedural learning: Evidence from cerebellar damage. Neuropsychologia 45:1208–1214.

Tronson NC, Taylor JR (2007) Molecular mechanisms of memory reconsolidation. Nature Reviews Neuroscience 8:262–275.

Tulving E (1972) Episodic and semantic memory. In: Organization of Memory, Tulving E, Donaldson EW (eds). New York, Academic.

Tulving E (1984) Multiple learning and memory systems. In: Psychology in the 1990's, Lagerspetz KMJ, Niemi P (eds). Holland, Elsevier.

Tulving E (1985) How many memory systems are there? American Psychologist 40:385–398.

Vertes RP, Eastman KE (2000) The case against memory consolidation in REM sleep. Behavioral and Brain Sciences 23:867–876.

Warrington K, Shallice T (1969) The selective impairment of auditory verbal short-term memory. Brain 92:885–896.

Warrington EK, Logue V, Pratt RTC (1971) The anatomical localisation of selective impairment of auditory verbal short-term memory. Neuropsychologia 9:377–387.

Wickelgren WA (1968) Sparing of short-term memory in an amnesic patient: Implications for strength theory of memory. Neuropsychologia 6:235–244.

Whalen PJ (1998) Fear, vigilance, and ambiguity: Initial neuroimaging studies of the human amygdala. Current Directions in Psychological Science 7:177–188.

Whittington M, Traub R (2003) Inhibitory interneurons and network oscillations in vitro. Trends in Neuroscience 26:676–682.

Zilli EA, Hasselmo ME (2007) Modeling the role of working memory and episodic memory in behavioral tasks. Hippocampus 18:193–209.