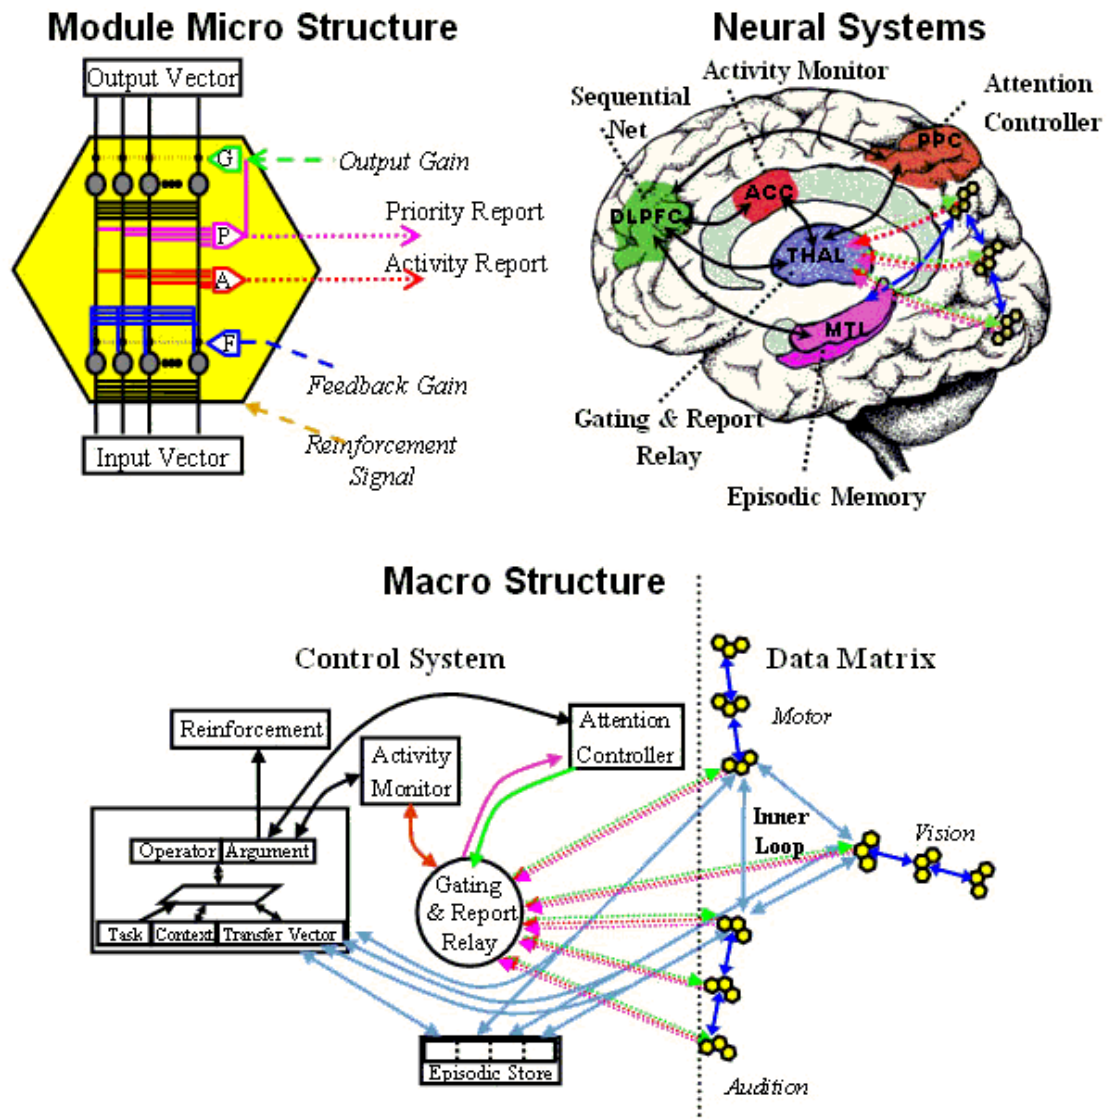


CAP 2 Model



Controlled & Automatic Processing: Behavior, Theory and Biological Mechanisms

Walter Schneider & Jason M. Chein

To appear in Cognitive Science 2003

Send Correspondence to Walter Schneider at wws@pitt.edu

[Abstract](#)

[Introduction](#)

[Definition of Automatic & Controlled Processing](#)

<u>Seven Empirical Phenomena of Automatic and Controlled Processing</u>
<u>Advantages of Two Different Processing Modes</u>
<u>Modeling of Automatic and Controlled Processing</u>
<u>Data Matrix of Modules</u>
<u>Simulation of Automatic Processing</u>
<u>Control System</u>
<u>Operations of Controlled Processing</u>
<u>Controlled Process Monitoring of Inner Loop Messages</u>
<u>Simulation of the transition from Controlled to Automatic Processing</u>
<u>Account of Seven Phenomena of Dual Processing</u>
<u>Functional Anatomy of Controlled & automatic processing</u>
<u>Relationship to Biological and Computational Models</u>
<u>Conclusion</u>
<u>References</u>
<u>Table 1: Summary of CAP2 Components and Features</u>

Figures

- [Figure 1.](#) Visual Search Paradigm
- [Figure 2.](#) Microstructure of a CAP2 module
- [Figure 3.](#) Macro-Structure of CAP2
- [Figure 4.](#) Priority learning simulation
- [Figure 5.](#) FMRI image of Control System
- [Figure 6.](#) Mapping of CAP2 Brain Structures

Abstract

This paper provides an overview of developments in a dual-processing theory of automatic and controlled processing that began with the empirical and theoretical work described by Schneider & Shiffrin (1977) and Shiffrin and Schneider (1977) over a quarter century ago. A review of relevant empirical findings suggests that there is a set of core behavioral phenomena reflecting differences between controlled and automatic processing that must be addressed by a successful theory. These phenomena relate to: consistency in training, serial versus parallel processing, level of effort, robustness to stressors, degree of control, effects on long term memory, and priority encoding. We detail a computational model of controlled processing, CAP2, that accounts for these phenomenon as emergent properties of an underlying hybrid computational architecture. The model employs a large network of distributed data modules that can categorize, buffer, associate, and prioritize information. Each module is a connectionist network with input and output layers, and each module communicates with a central Control System by outputting priority and activity report signals, and by receiving control signals. The Control System is composed of five processors including a Sequential Net, an Attention Controller, an Activity Monitor, an Episodic Store, and a Gating and Report Relay. The transition from controlled to automatic processing occurs in this model as the data modules become capable of transmitting their output without mediation by the Control System. We describe recent progress in mapping the components of this model onto specific neuroanatomical substrates, briefly discuss the potential for applying

functional neuroimaging techniques to test the model's predictions, and its relation to other models.

Acknowledgements: This work was supported in part by grants from the Office of Naval Research (ONR010360) and the National Science Foundation (NSF9873465)

The view that human cognition may comprise two different types of processing, controlled and automatic, has been a theme in the psychology literature for over a century (e.g., James, 1890). This notion of dual processing has been prominent in the work of Richard Shiffrin and his students for the past thirty years. Shiffrin's early work with Atkinson (1968) detailed the role of controlled processing in studies of short-term memory and verbal learning. In the early 1970s, he and his students (Shiffrin & Gardner, 1972; Shiffrin, McKay, & Shaffer, 1976) performed attention studies indicating that multiple channels could be processed in parallel, a result that ran counter to much of the attention research of the time. Building on these earlier findings, Schneider and Shiffrin (1977) and Shiffrin and Schneider (1977) published a set of companion papers that set out the theoretical and empirical basis for much of the work on automaticity that has emerged in the ensuing decades (as illustrated by the over 4800 citations of this work). There has followed an extended effort to develop an empirical and theoretical understanding of automatic and controlled processing (Anderson, 1992; Logan, 1980; Pashler, Johnston, & Ruthruff, 2001; Stanovich, 1987).

The present paper provides a summary and perspective of this evolving work, and considers some of the biological mechanisms that may underlie dual processing. While our early work dealt with issues of visual search and basic executive function, subsequent research has elaborated on the role of controlled processing by applying the basic framework to understanding novice rule-based processing, problem solving, fast learning, and consciousness. These applications of the theory have been tested in simulations with a computational model (CAP2) designed to capture key aspects of the dual processing theory. In other research, we have begun to explore the physiological mechanisms mediating these different forms of processing (Schneider, 1999).

The mapping between empirical findings, theory, and biological mechanisms is complex. This complexity may be explained, in part, by recognizing that human performance often results from an interplay between automatic and controlled processing and these processes may involve different cortical mechanisms. This interplay is mediated by systems that have evolved to satisfy the need for operation in a complex environment, wherein attention must be guided to selectively process critical stimuli. From an evolutionary perspective, gradual increases in the complexity of behaviors supported in the human repertoire may have seeded the development of a precursor executive control system that later evolved to support fast (e.g., single-trial) learning mechanisms, and the ability to override learned associations, that are characteristic of human cognitive performance.

Definition of Automatic & Controlled Processing

The basic nature of automatic and controlled processing was laid out in our earlier papers. In Schneider & Shiffrin (1977), an *automatic process* was defined as the activation of a sequence of nodes that “nearly always becomes active in response to a particular input configuration,” and that “is activated automatically without the necessity for active control or attention by the subject” (p. 2). This ability for a process to occur in the absence of control and attention by the subject is perhaps the most salient feature of an automatic process, and was the basis for referring to such processing as “automatic”. In general, automatic processes “operate through a relatively permanent set of associative connections ... and require an appreciable amount of consistent training to develop fully”

(Schneider & Shiffrin, 1977, p. 2). An **automatic attention response** is a special type of automatic process that directs attention automatically to a target stimulus (Schneider & Shiffrin, 1977).

In contrast to automatic processes, Schneider and Shiffrin (1977, pp. 2-3) defined a **controlled process** as “a temporary sequence of nodes activated under control of, and through attention by, the subject.” Furthermore, controlled processes are “tightly capacity limited, but the costs of this capacity limitation are balanced by the benefits deriving from the ease with which such processes may be set up, altered, and applied in novel situations for which automatic sequences have never been learned.”

The contrast between automatic and controlled processing was initially studied using extended **consistent mapping** (CM) training. A consistent mapping task is one in which the response to the stimulus is consistent across extended periods of time (e.g., in a search task, the set of target stimuli is constant throughout the experiment). Under consistent mapping, automatic processes can develop slowly as repeated stimuli are attended to. Although there may be marked performance improvements even within the first few trials (Logan, 1992), full automaticity typically requires hundreds of trials to develop. In **varied mapping** (VM) training, the relationship of the stimulus to response mapping varies from trial to trial (e.g., in a search task, a stimulus that is assigned a given response on one trial is assigned a different response on the next trial). With varied mapping, the prior and current associations are incompatible, thereby precluding automaticity and the development of an automatic attention response.

Seven Empirical Phenomena of Automatic and Controlled Processing

There has been a long and rich empirical history in the development of dual processing theory, as detailed in a review by Shiffrin (1988). Additional recent reviews include Näätänen (1992), Bargh (1992), and Schneider et al. (1994). In this section we will explain the core phenomena of controlled and automatic processing for the purpose of later relating them to our theory and toward understanding the biology of the two processing modes. There are seven behavioral phenomena that should be explained by a comprehensive theory of automatic and controlled processing.

(Insert Figure 1 about here)

The major empirical paradigm in which these phenomena are explored is the **search paradigm**. Figure 1 illustrates the search paradigm and the consistency manipulation. The subject is given a memory set of typically one to five items that they must remember. Then the subject sees a set of sequentially (single-channel) or simultaneously (multiple-channel) presented **probe** items that he or she must search through. If any of the probe items matches any of the items in memory, a **target** is detected, and the subject makes a positive response. When a probe item does not match any of the memory set items, it is referred to as a **distractor**, and the subject makes a negative response or no response at all. For example, the subject may be given a **memory set** of the letters “KJTL” to remember, and then a probe (e.g., “L” or “P”). The subject compares each memory set item to the probe item, looking for a match. In VM, items can be a target on one trial, but then become a distractor on the next trial. Therefore, if the task has a varied mapping, then a given stimulus may require varied responses from trial to trial. In contrast, in CM, items that have been targets on one trial never become distractors in subsequent trials.

Therefore, if the task has a consistent mapping, then a given stimulus always requires the same response¹.

The first relevant phenomenon observed in this paradigm is that **extended consistent training is required in order to develop automatic processing, while controlled processes can be established in a few trials and under varied mapping situations**. As mentioned above, automatic search training develops after extended consistent search. For automatic search in detection experiments, the rate and effectiveness of learning are a function of the degree of consistency (Fisk & Schneider, 1982). Shiffrin & Schneider (1977, experiment 3) showed that it is specifically the consistent mapping, rather than learning to categorize items rapidly (category search), that is critical for automaticity. Subjects began the experiment under varied mapping conditions in which they learned that letters belonged to one of two categories (GMFP vs. CNHD), but the category being searched for was varied from trial to trial (e.g., on one trial they searched for the letters GMFP with distractors CNHD, and on the next trial they searched for the opposite set of CNHD with distractors GMFP). The category training continued for 25 sessions until the memory set size effect for searching for 2 versus 4 members of the category was the same, and performance was at asymptote. At this point, the search was made consistent, with one category being always the target. There was a dramatic improvement in processing speed, and a reduction in effort at the point where consistent search was introduced. This result supports the view that automaticity, and the automatic attention response, emerges only when the stimulus-response mappings remain consistent.

It is important to further note that the consistency of training need not be at the exemplar level. Fisk & Schneider (1983) showed that, in category search, repetition of distinct items belonging to a common category, and not just repetition of the specific exemplars, promotes search improvement. They also demonstrated that CM learning of a category transfers its performance advantages, to a large extent, to other members of the category even on first presentation of the new members.

The second relevant phenomenon is that **automatic processing is fast and parallel, while controlled processing is slow and serial**. In visual search tasks, for example, when using a VM paradigm, search is slow and serial. Under VM conditions, search time is typically 40-60 ms per letter when a set of letters are the targets (Johnsen & Briggs, 1973; Schneider & Shiffrin, 1977) and 200 ms per category comparison when items must be detected as members of a semantic category (Fisk & Schneider, 1983). The graph in Figure 1 illustrates the difference in response times between VM and CM search. Indicative of the serial nature of controlled processing is the finding that VM search times increase linearly with the number of comparisons that must be made, both in memory (memory set) and in the display (search set). The persistence of serial processing in VM search is also suggested by examination of the variance associated with search reaction times. Specifically, for both positive (when the target is found) and negative (when the target is not found) response trials, variance increases as the number of required comparisons is increased. This pattern of mean and variance increases for positive and negative responses provides strong evidence for serial search (Schneider & Shiffrin 1977,

¹ The search task can be made more challenging by presenting multiple probe stimuli simultaneously in the display (e.g., 2 to 4 letters). Such multiple-channel probe displays can contain both targets and distractors. Consequently, while the response to targets remains consistent, the response to distractors will vary depending on whether they co-occur with a target.

Appendix F; Townsend & Ashby, 1983). These findings contrast with those from CM search tasks, which show non-linear parallel search patterns with fast slopes (e.g., a category search slope of 2 ms per category). As an important aside, it should be noted that controlled processing search is typically serial across memory set items, across items in a multiple-channel display, and across stages (e.g., perception and motor response). However, controlled processing can be parallel across items in a multiple-channel display when there is a single distinguishing feature being searched for (e.g., Treisman & Souther, 1985) or when the memory set size is one item (Schneider & Shiffrin 1977, appendix J).

A third important phenomenon is that **automatic search requires little effort and can operate in high workload situations, whereas controlled processing requires substantial effort and interferes with other controlled processing tasks.** This difference is very salient in subjective reports, in which subjects typically indicate that CM search does not feel mentally taxing, while VM search requires continuous effort. In dual task situations, it has been demonstrated that CM search can occur in parallel with other CM or VM search tasks, whereas VM search trades off with other VM search tasks (Schneider & Fisk, 1982).

A fourth phenomenon is that **automatic processing is rather robust to stressors.** In studies of the effects of alcohol, fatigue, stress, and vigilance on performance (Fisk & Schneider 1981, 1982; Hancock, 1986; Heuer, Spijkers, Kiesswetter, & Schmidtke, 1998), performance under automatic processing is far less sensitive to such stressors than is performance under controlled processing.

A fifth phenomena (one that at the time we demonstrated it was so striking that it influenced the naming of “controlled” processing in our theory) is the **difference in cognitive control that can be applied to automatic and controlled processes.** Specifically, once a process becomes automatic, it becomes difficult to control. The prototypical example of this phenomenon is the difficulty of control exhibited in the Stroop (1935) task. Likewise, Shiffrin and Schneider (1977, experiment 1) showed that subjects took three times longer to unlearn and relearn an automatic search set than it took to learn the search set originally. CM foils, stimuli that had been previously trained as CM targets but were now meant to be ignored, demanded attention (i.e., produced the automatic attention response) even when they occurred in locations known to never contain a target (Shiffrin & Schneider 1977, experiment 4d). These effects could continue for thousands of trials. In contrast, controlled VM search could be altered in a single trial, allowing for variation in what was being searched for, what locations the items were in, and of the nature of the required response, without producing significant changes in performance level.

A sixth phenomenon is that the degree of **learning is dependent on the amount and type of controlled processing, while there is little learning in pure automatic processing.** Fisk & Schneider (1984) showed that in a VM search task, involving controlled processing, the ability to recall having seen a distractor was a function of the number of times a distractor was compared (and also of the type of task, word versus category search). In contrast, in a CM search task, involving automatic processing, subjects can repeatedly categorize words into a semantic category with little, if any, memory modification (learning) resulting. The strong link between learning and

controlled processing, and the lack of such a link in automatic processing, is an important qualitative separation of these processing modes.

The seventh phenomenon is that the **automatic attention response is dependent on the priority assigned to a stimulus itself, rather than on the context in which the stimulus occurs.** An automatic attention response could in principle be triggered by the simple presence of a stimulus, or by the presence of the stimulus only when it appears in a particular context (e.g., when it is deemed to be a target, when it is held in working memory). Empirical results support the hypothesis that the automatic attention response results from a priority code based solely on the stimulus input. That is, every stimulus has an intrinsic (and modifiable) probability of attracting attention. If the probability is high, then the stimulus is likely to grab attention without requiring the intervention of a controlled process (i.e., is likely to evoke an automatic attention response). When there are several stimuli with high priority codes, the one with the highest priority wins out. These principles account for the nature of skill transfer and for the specific effects of CM training. For example, Shiffrin, Dumais, and Schneider (1981) examined training and transfer in visual search tasks. In these experiments, subjects were trained to search for particular sets of letters while ignoring other sets of letters. For example, subjects may have searched for a given set of target letters, set A, in a set of distractors, set B (we refer to this as A[B]). Training with A[B] shows positive transfer to scenarios where subjects must search for previously trained targets in a new set of distractors (A[N], where N is a new set of letters) and where subjects must search for new targets in old distractors (N[B]). Presumably, in each of these conditions, prior experience with searching for stimuli in set A leads to an increase in their attentional priority, which then transfers across contexts, while experience ignoring stimuli in set B leads to a decrease in their attentional priority to a level below that of novel stimuli, which also transfers across contexts. Accordingly, negative transfer occurs when old targets are to be ignored (N[A]), old distractors are to be detected (B[N]), and when the targets and distractor sets are reversed (B[A]). Interestingly, subjects succeed in developing automaticity when search sets have strict priority A[B] and B[C]. However, completing the “triangle” by adding search for C[A] violates strict priority and causes one of the search conditions to show VM search behavioral patterns. We have also attempted a wide range of manipulations of internal state (e.g., memory) and external state (e.g., background of the display) and found that priority code changes are independent of context (e.g., searching for A[B] on a green background and B[A] on a red background does not allow for the development of automatic attention responses (Schneider, unpublished). Interestingly, this type of priority learning shows location specificity, with a fall off in transfer to more distant spatial locations.

Advantages of Two Different Processing Modes

Dual processing mechanisms would likely not have evolved unless there were survival advantages to having both modes of processing. We assume that automatic and controlled processing are two qualitatively different forms of processing that provide complementary benefits. Our simulations suggest that a single process alone can not provide both the fast learning of controlled processing and the high speed parallel robust processing of automatic processing. So, although it may be less parsimonious to assume two different modes of processing, we argue that there are sufficient survival advantages

to a two-process system over a unitary architecture to have allowed a dual-process system to evolve. The survival advantages to having both controlled and automatic processing are analogous to the non-overlapping and overlapping benefits of having rod and cone vision. With controlled processing: 1) the fundamentals of new skills can be acquired quickly (e.g., one trial learning to escape when a life threatening stimulus appears), 2) critical stimuli can be attended while ignoring normally relevant stimuli (attend to a child in a crosswalk while inhibiting the prepotent response to accelerate on a green light) 3) variable bindings that allow general operations to be applied to temporarily relevant stimuli can take place (e.g., after eating a novel food, search for it in the environment), 4) learning can be passed between individuals by instruction or observation (rather than shaping), and 5) goal directed behavior can be planned and executed. However, due to the slow execution, high effort, and poor robustness of controlled processing, it can operate on only a small number of stimuli at any time, and any skill acquired during controlled performance may not be sufficiently robust to resist rapid decay or deterioration in the presence of stressors. Further, if a task requires the coordination of many sensory/motor inputs, the slow, resource limited nature of controlled processing will be a serious limitation (e.g., imagine trying to ski down a difficult slope using verbal rules to plan and execute motor movements). Despite taking a long time to acquire, automatic processing has the advantages of being robust under stress, leading to long-term retention of associated skills, and allowing many processes to occur in parallel.

Modeling of Automatic and Controlled Processing

The modeling of automatic and controlled processing has been an active forum in cognitive science over the past twenty-five years. The initial Schneider and Shiffrin (1977) models provided a quantitative account of the behavioral data of controlled processing search accuracy and reaction time effects. Several other approaches to modeling controlled and automatic processing have since been demonstrated. For example, production system based computer modeling approaches (ACT-R, CAPS, SOAR) have been applied to a wide range of cognitive tasks in which automaticity may develop (particularly in problem solving behavior; Anderson, 1992; Just, Carpenter, & Varma, 1999; Young & Lewis, 1999). Connectionist and parallel distributed processing models (McClelland & Rumelhart, 1988) have also been applied to address the development of automatic processes. Schneider & Detweiler (1987) described a connectionist control architecture (CAP2) to account for the role of controlled processes in working memory phenomena.

In this section we provide a review and extension of the Controlled and Automatic Processing 2 (CAP2) architecture (see also Schneider, 1999) that we have used to explore dual processing theory. The CAP2 model seeks to account for the phenomena of automatic processing, and to explain the role of controlled processing in a wide range of cognitive tasks.

Conceptually, CAP2 is a hybrid cognitive architecture, incorporating both symbolic (e.g., ACT-R) and connectionist (e.g., PDP) elements. In actuality, the CAP2 architecture is implemented with entirely connectionist components, but has networks that operate as sequential control structures that can configure the network to behave as a production system (e.g., to implement condition-action rules with variable binding based operations). The hybrid architecture is intended to capitalize on the complementary

strengths exhibited by the two modeling approaches, as they relate to controlled and automatic processing. Specifically, production system architectures account well for symbolic variable binding behaviors, in which an algorithm in memory takes on a specific, yet arbitrary, input as the value of one of its variables (e.g., in VM search, an incoming stimulus is mapped to a variable representing the trial input, and then the particular form of the stimulus is compared to internal representations of the target stimuli). Symbolic architectures can easily perform such variable binding tasks for arbitrary symbols (e.g., determine whether the following sequence is a palindrome of letters, “tzppzt”), while it is very difficult to build connectionist architectures that perform such tasks without extended training of the stimuli (Touretzky, 2002).

Symbolic processing has the disadvantage of producing representations that are often brittle and exhibit limited transfer. Connectionist processes, in contrast, show good generalization and parallel processing, but often at the cost of very long learning times (tens of thousands of trials), and an inability to transfer to analogous tasks that do not include similar stimulus elements.

These properties of symbolic and connectionist architectures can be mapped onto the relative advantages of controlled and automatic processing. Recall that automatic processes operate through a relatively permanent set of associative connections and require an appreciable amount of consistent training to develop fully. Automatic processes therefore behave much like most connectionist networks (after they have been trained). In contrast, a controlled process “may be set up, altered, and applied in novel situations for which automatic sequences have never been learned” (Schneider & Shiffrin, 1977, pp. 2-3). Accordingly, controlled processing proceeds similarly to processing in most production system architectures.

Table 1 provides a summary description of the CAP2 model architecture. This model enables predictions at the micro level (type of units), macro level (organizational structure of cortical connections), and process levels (nature of executive function). Providing a biologically feasible model for human cognition requires a detailing of a comprehensive architecture. It is a challenge to relate the multiple modes of human cognition to the rich complexity of brain structure. Most models based on neural principles seek to demonstrate the ability of parsimonious architectural elements to predict human behavioral performance (e.g., computation by populations of additive connectionist units, as in McClelland & Rumelhart, 1988). In contrast, the **modules** in CAP2 attempt to capture the computational richness of the diverse neuronal assemblies that comprise cortical modular columns (hyper-columns), which are found to recur throughout the cortex with regionally specialized connection patterns (White, 1989; Martin, 1988; Felleman & Van Essen, 1991; Goodman et al., 2001).

(Insert Table 1 about here)

One of our goals in developing the CAP2 model has been to implement symbolic and connectionist processing through the use of a modular connectionist architecture and to provide explicit quantitative representations of both the development of automaticity and the nature of controlled processing. The model achieves this goal by employing a **Control System** that monitors and modulates activity in a large **Data Matrix** of connectionist modules. The connectionist modules in CAP2 are intended to mimic the modular quality of cortical structure. As prefaced above, the brain is composed of many modules, each containing a large number of cells (some 10,000 hyper-columns, each

containing roughly 40,000 cells in the human brain; see Martin, 1988). Each of the modules in CAP2 can receive input vectors, and can locally buffer, categorize, prioritize and associate (learn) an output vector that can be sent to other modules. The Control System is implemented as a set of interlinked connectionist processors (with properties resembling a symbolic architecture) that monitor and set control signals sent to the network of connectionist modules. These control signals alter functional connectivity within the architecture, thereby enabling a wide variety of cognitive operations.

(Insert Figure 2 about here)

Data Matrix of Modules. The module microstructure (Figure 2) involves two layers of units. An input vector of activity excites the input-layer of units, which have excitatory auto-associative connections (i.e., project back onto themselves), as well as connections to a set of output units. The auto-associative connections clean up and categorize the incoming signal (see Anderson, J. A., 1983). The output-layer of units send a vector of activation to other modules. The module architecture can be related to the classic connectionist three layer structure (input layer, hidden layer, output layer) where the input layer maps to the CAP2 module input layer, the hidden layer maps to the CAP2 module output layer, and the output layer maps to the CAP2 input layer in the module in the next higher tier of the hierarchy. The output layer is gain modulated such that output vectors are transmitted only when the output gain is high.

There are three module input control signals and two module output report signals. Providing control signals to a module allows the symbolic operations of the Control System to influence their behavior. The first input control signal is the regional **feedback gain**, which modulates the auto-associative feedback on units in the input layer. This feedback influences the coarseness of categorization of the input (Anderson, J. A., 1983), and also affects whether an input vector will be buffered or permitted to decay (Shedden & Schneider, 1990). Modulating the input feedback allows the network to dynamically alter the categorization criterion (e.g., when an ambiguous input occurs, the input-feedback can increase until there is a categorization to a previously learned state), and to alter the extent of exogenous versus endogenous activation of a module (i.e., the extent to which the module is receptive to an arriving input vector). The second control signal is the **output gain**, which determines attentional selection. The output gain of a module is a scalar multiplication of all the output units. The output gain signal can increase or decrease a module's output gain to influence the transmission of the module's output vector. This signal is the principle mechanism of attentional control (attending to a module is equivalent to temporarily increasing its output gain so that it can transmit). The third control signal is the global **reinforcement signal**. This signal causes a change in the connection weights within and between modules, and is therefore the triggering event that causes learning.

Learning within a module actually takes two forms; associative learning and priority learning. That is, when the reinforcement signal is given, any module that recently received an input vector and transmitted an output vector will alter its connection strengths in two ways. First, **input-output associative learning** changes connection weights between the input vector and the output vector (through connectionist learning rules such as back-propagation learning, see McClelland & Rumelhart, 1988) so that the input vector will ultimately come to evoke the appropriate output vector without external

influences. Second, **priority learning** changes the connection weights between the input vector and the priority report (see below) to code the importance of the message.

There are two module scalar output report signals, the **activity report** and the **priority report**, that provide the Control System with information needed to guide module selection and to detect memory matches (i.e., recognition). The report signals are critical for keeping the Control System from falling into the “homunculus trap” of being overwhelmed by overly complex input. The modules do the work of processing a vector of input (e.g., activation of 10,000 units) and represent the data to the Control System with just two simple scalars representing activity and priority. A basic task of the Control System is to simply attend to the most important message, based on the module priority and activity reports. One can think of the module’s scalar coding as a representation of the deemed importance of the message for transmission (e.g., like a person raising their hand to a particular height as an indication of how important they think their contribution to a discussion might be). The “executive” Control System can then be implemented as a simple winner take all network that identifies the “active” (above a threshold) module with the highest priority signal (as explained below). With module reports, a simple executive can thus effectively manage communications in the network.

The activity report, provides a metric of activity in the module (i.e., the sum of the activity across all the units active in the input layer of a module; see Shedden & Schneider, 1990). This report provides a metric of the degree of match (correlation) between multiple simultaneous input vectors to the module (Schneider & Oliver, 1991). The controller can limit search to those channels that have high activity reports, and can skip blank channels (e.g., attend to the ear that is receiving a message and ignore the other ear), thereby accounting for the lack of load associated with having to monitor potential, but empty, channels (Shiffrin & Schneider, 1977, experiment 4a).

The second module output report, the **priority report**, determines which modules have the most important message to transmit. The priority report is the result of an association of the input vector to the priority output unit for that module. Important stimuli (e.g., one’s name) develop a high priority and are likely to be attended to. The priority report has a distant and a local effect. The distant output goes to the Control System, indicating the importance of the message. The local signal connects to the output gating units of the module, and enables a brief output of the vector for high priority signals (even without an output gain signal from the Control System). This local connection is the basis for automatic output from the module.

(Insert Figure 3 about here)

The macro level structure of CAP2 involves a **Data Matrix** of the above modules and a separate **Control System** that modulates communications within the system. Figure 3 illustrates the model’s macrostructure with the Control System on the left, and the Data Matrix of modules on the right. The structure of the Data Matrix of modules is illustrated on the right of Figure 3. Each region (e.g., vision, audition, motor) is a hierarchy of tiers, with each tier composed of hundreds of modules. There may be many tiers in a region, and the modules within a tier can connect to multiple modules above and below them. Tiers of the model correspond with hierarchical processing levels in the brain. In vision, for example, there are an estimated 32 specialized cortical processing areas (e.g., V1, V2, V3, V3A, V4, TEO, etc.; see Baizer, Desimone, & Ungerleider, 1993; Felleman & Van Essen, 1991), and each area would be represented as a tier in the

model.² Across the brain, there are a substantial number of such tiers (Worden & Schneider, 1995, estimate that there are 500-1000 cortical areas based on differentiation of primate cortex) with a modest number of tiers for a given function (e.g., object identification involves twenty tiers).

The most peripheral tier of each region corresponds to primary cortical areas. Sensory peripheral tier modules receive input from the sensory receptor systems. Meanwhile, motor peripheral tier modules output to motor systems.

The top tier of each region is interconnected with other top tier modules on a circuit referred to as the **inner loop** (Schneider & Detweiler, 1987; see right side of Figure 3). Given their hierarchical position, these inner loop modules are likely to be found in association cortices. The modules on the inner loop have a special status in three respects. First, the inner loop of modules allows cross-region transmissions (e.g., vision to motor output). One important function of the inner loop is therefore that it provides cross modality communication while still allowing parallel computation within regions to occur without interference. For example, visual processing and motor processing can occur in parallel because they do not share connections (and hence potential for communication interference) except at the top tier. Second, the inner loop modules communicate vector messages to the Control System. Consequently, modules on the inner loop are special in that their output vectors are the only representations (non-scalar information) available for inspection by the Control System. This pathway for communication allows the productions established by the Control System to incorporate vector information represented in the inner loop modules, but not in the lower tier modules (see p 8 below). Third, the inner loop modules project to an **episodic store** repository that supports the storage and retrieval of information on the inner loop, thus enabling previous inner loop states to be “reloaded”.

Simulation of automatic processing. Automatic processing in a CAP2 simulation is a direct result of having a high priority report producing an automatic transmission of a module’s output vector in the absence of controlled processing input. More specifically, automatic processing occurs when priority coding is sufficiently high that it can trigger transmission, without waiting for an output gain signal to be input from the Control System. Gupta and Schneider (1991) examined CM and VM learning in a simulation of a search task. In simulated CM trials, the priority report increased the differentiation between eight target and eight distractor stimuli over 400 trials. The result was that the reaction time functions for two to four item searches became flat after 200 trials, consistent with prior behavioral data (Kristofferson, 1972; Schneider & Shiffrin, 1977, experiment 2). In VM search, there was no target/distractor differentiation, and hence search remained serial throughout training (as in Kristofferson, 1972; Schneider & Shiffrin, 1977, Experiment 2).

The CAP2 model provides a computational account of automatic and controlled processing that is consistent with the earlier Schneider & Shiffrin (1977) description. Recall that we defined an automatic process as a sequence of nodes that “(nearly) always becomes active in response to a particular input configuration ... without the necessity of active control or attention by the subject” (Schneider & Shiffrin, 1977, p. 2). In CAP2, after CM training, a stimulus that evokes a high priority response will output the vector to

² While tiers are organized in the current model as a strict hierarchy, we recognize that some cortical processing streams have convergent and divergent pathways (e.g., V2 to both V4 and MT)

other modules without the need for controlled processing, thereby demonstrating automaticity.

It should be noted that this does not imply that attention cannot influence processing.³ The Stroop (1935) task illustrates the interaction of automatic and controlled processing. The transmitting process may be automatic in the sense that, for word reading, a short automatic transmission occurs to a word that is attended. The word stimulus is then transmitted to higher states that would ultimately produce the wrong response in an incongruent Stroop trial (e.g., a trial in which the word “RED” is printed in green ink, and the subject is to respond by saying “green”). During Stroop training, the automatic word transmission would occur for a short time (less than half a second), and would shut down. Thereafter, vectors from modules attending to the ink color could be transmitted to the response system. Early in each trial, the motor system must be inhibited (kept at high feedback and low gain to block motor transmission of the word), and then, after the automatic transmission of the word drops down, the motor output is allowed to respond with the transmission of the ink color.

Control System. The Control System in CAP2 is designed to dynamically reconfigure the network to perform different tasks by monitoring the report signals (activity and priority), and then transmitting control signals (gain, feedback, and reinforcement) back to the Data Matrix of modules. The left side of Figure 3 illustrates the internal structure of the Control System as implemented in CAP2.

There are five processors in the CAP2 Control System architecture, each with an information-processing role and likely cortical substrate. The **Sequential Net (SN)** performs sequential control operations altering information flow in the Data Matrix. The SN is a connectionist implementation of a production system, and performs condition-action rules based on the contents of the modules and the task vector. In a visual search task, the SN sequentially activates the modules containing the memory set and the modules containing and encoding the probe letters, monitors the degree of activity/match of the inputs, and releases a response vector when there is a match (further detailing of the SN is provided below).

The **Attention Controller** tracks the module priority signals from the Data Matrix and changes output gains based on both top-down commands from the SN and bottom-up activity reports (e.g., the SN requests a scan of the visual scene from left to right, the bottom-up reports identify three sets of modules of high priority, the Attention Controller sequentially sends a high output gain signal to each of the high priority modules). The Attention Controller is more than a relay station since it performs actions such as correcting the addressing of modules with changes in head and eye position, tracking positions that have been visited to minimize reprocessing of recently visited locations, and determining the rate at which items are scanned.

The **Activity Monitor** processes the activity reports arriving from the Data Matrix, sets match thresholds for these reports, and communicates information regarding supra-threshold activity on the network to the SN. Like the Attention Controller, the Activity Monitor can constrain its scope to modules specified by top-down commands from the SN (e.g., monitor activity at the visual shape level) and may be also be driven by bottom-

³ If the controlled output and automatic output increase the output vector gain, controlled processing can influence automatic processing. If controlled processing reduces the gain of a later stage or increases the feedback of a stage, the degree of activation of the later stage can be influenced.

up signaling based on the activity reports of the Data Matrix. Such activity reports from the Data Matrix reflect the degree of match for comparisons made in a given module, and also signals the decay of traces in a module. The Activity Monitor locally adjusts match thresholds to be used in a task, and buffers the magnitude of supra-threshold matches along with the address of the reporting module. This buffering allows the SN to check the match outcomes intermittently.

The **Episodic Store** records the contents of vectors in the inner loop Data Matrix modules (association cortex) and also of the SN output vector, thus allowing the Control System to “recall” states of the inner loop modules in order to initiate plans.

The **Gating & Report Relay** is a central routing nucleus for report and gain control signals. This processor performs two functions: 1) to relay output gain control signals, emitted from the Attention Controller, to their target modules, and 2) to communicate a set of high priority and activity reports to the Attention Controller and Activity Monitor, respectively. Thus, report signals from thousands of modules are compressed and prioritized into vector messages.

The SN is a high level central executive system. It is “executive” in that it works on preprocessed data from the Data Matrix. In other words, the SN does not process information vectors occurring throughout the Data Matrix, but rather operates on preprocessed reports from the data modules (and on vector transmissions from the inner loop, as explained below). Again, by having the controller operate on preprocessed signals, the notion of homunculus that must process all information in the system is avoided, but the executive consequently has no access to underlying codes (unless they get to the inner loop).

Schneider and Oliver (1991) implemented SN controlled processing as a connectionist sequential recurrent network (adapted from Elman, 1990). The controlled processing network differs from the modules of the Data Matrix both in its inputs/outputs, and in its function. This SN has as input the report signals of the modules that have been preprocessed in the Activity Monitor. It also has three additional sources of input (see Figure 3): message vector inputs from the inner loop of modules, a context vector that represents the current environment (Schneider & Detweiler, 1987), and a task vector that represents the task demands (Schneider & Oliver, 1991). The SN can also output vectors to the inner loop, and thereby influence information available to the Data Matrix. This control network can implement a wide range of functions by loading different task vector programs that alter the way in which information can flow in the Data Matrix.

The base **operations of controlled processing** include: 1) altering the gain of specific modules (attention); 2) attending to high priority messages; 3) changing feedback (categorization, buffering, and clearing of a module message vector); 4) comparing vectors; 5) sending reinforcement signals; 6) sending an output vector to the inner loop; 7) configuring the Data Matrix (setting up a dynamic set of transmitting modules); 8) binding new associations; 9) performing goal based executive operations; 10) using context memory to return to previous tasks; and 11) recalling information from declarative memory (see Schneider, 1999).

Again, controlled processing can be implemented as a sequential recurrent network that performs serial operations of a task, such as those required when solving logic problems or performing visual search (before training). Although controlled processing

is implemented using connectionist units, it performs operations like a production system (see also Lebiere & Anderson, 1993). Since data modules can contain vectors representing a particular class (e.g., lexical items) and controlled processing can execute condition-action rules, production system operations can be implemented. The “condition” part of the productions can reference information in the control processors and from vectors in the inner loop modules (i.e., in top-tier modules).

In simulation studies, Schneider and Oliver (1991) demonstrated that the SN could learn and execute multi-step programs, and could implement simple finite state grammars (see also Elman, 1990; Cleeremans et. al, 1989). In a simulation of controlled processing, the SN learned six concurrent sequential programs of up to 25 steps each to perform digital logic problem solving for six gate types (AND, OR, XOR, NAND, NOR, XNOR) in 120 total trials (humans required 216 trials to reach accurate performance). The control network “learned the programs” through direct training of the vectors. That is, the task vector was set, then for each step of the program the output/input states of the network were set and the SN learned the mapping between the task state, the sequential steps of the program, and the final state using back-propagation to change the connection weights. The sequences did include conditional states (e.g., if the inputs were all ones, all zeros, or mixed) and each conditional task was rehearsed on separate trials depending on the random sequencing of the trial conditions. This is very similar to what happened in human subjects, in that humans subjects were given explicit verbal instruction (e.g., attend to the inputs of the gate: if all ones then code as all ones, if all zeroes then code as all zero, otherwise code as mixed). It took the simulation an average of 20 repetitions per task program to learn the mapping. The rapid learning illustrates that using a SN that works with report and control signals to control a network of connectionist modules yields a connectionist executive system that can produce very fast learning of complex routines.

To implement the typical search task in controlled processing requires: 1) loading a set of lexical modules with the memory set, 2) monitoring the set of active visual channels, 3) sequentially attending (by increasing output gains) to each active set of modules, 4) monitoring the activity report of a visual shape module to determine if it is above threshold (see discussion of vector comparison below), 5) increasing the gain of an output module that is above threshold to produce a positive response, or shifting attention to the next remaining active display module if the prior module was not above threshold, 6) shifting attention to the next remaining memory module if the prior step did not produce an output, 7) making the negative response if no output has occurred, or sending a reinforcement signal if a positive response (output) occurred.

In CAP2, the Control System can trigger a reinforcement signal. This signal results in a change in those modules that recently transmitted, altering their associated priority code. If a positive reinforcement signal occurs, the mapping of a message vector within a module increases connections to the priority units to increase the priority report (Gupta & Schneider, 1991). If the module transmitted and there was no reinforcement signal, then priority would be reduced. Reinforcement also changes the within-module connection matrix so that the input comes to evoke the output.

An important aspect of the ability to associate vectors based on input-output relationships is that controlled processing can rapidly build arbitrary associations by loading vectors and altering the transmission. For example, in search tasks, it is typically

arbitrary what specific response is assigned to targets and distractors (e.g., push the “Z” key for targets, and the “/” key for distractors). When a target (e.g., “CAT”) is detected, controlled processing sets up vector transmission sequences (transmit “CAT” received, then release previously loaded output response “Z” key). Learning causes a local association of the “CAT” input to “Z” output in the module that had an input from the match and released the response. The Control System mediated association learning contrasts markedly with learning by “shaping,” which requires much more time to acquire the skill (e.g., a monkey often requires months of training for tasks that humans can acquire in minutes).

This ability of the Control System to stimulate associations between arbitrary vectors enables learning from instruction or imitation (see Schneider & Oliver, 1991; Schneider Pimm-Smith 1997), greatly speeding learning. In CAP2, instructing the Control System to execute a controlled processing program allows rapid learning of a new tasks and of “arbitrary” input-output patterns. CAP2 implements typical Hebbian learning such that a vector-in followed by a vector-out will result in changes in the connection matrix so that the input vector can ultimately evoke the output vector.

A very important role of controlled processing is to **speed the development of automatic processing by task division, and by learning from instruction**. In a simulation study (Schneider and Oliver, 1991), we compared connectionist back-propagation learning to learning in a connectionist network under the direction of a control processor (that could specify the levels and intermediate states for the task). For the six input logic gate task, back-propagation connectionist learning required 10,885 trials to reach criterion accuracy and reaction time. In contrast, the control processor assisted network required only 120 trials (91 times faster)⁴ to reach the same accuracy level. Still, performance remained slow and serial since it was dependent on the SN sequential operations. With continued training of the control processor assisted network, a shift to automatic processing (speeded reaction times and enabling independence from the Control System) was achieved after 948 trials (11 times faster than the back-propagation network). This learning speed-up was achieved by having the control processor divide the logic gate task into three subtasks (1- categorize input as all ones, all zeros, or mixed; 2- code gates as either AND, OR, XOR; and 3- perform negation when needed).

The ability to learn from observation or instruction through the transfer of human learning without laborious shaping represents a very important jump in cognitive function. Biederman & Siffrar (1987) remarked on an intriguing contrast in humans wherein learning can occur either through simple repetitive reinforcement, or through direct instruction. During the 1930s in the USA, the skill of identifying the sex of chicks by visual inspection was trained by having each student deduce the sex of the chick, and then having an instructor provide the correct answer. This training, which resembles the exemplar training used in connectionist modeling, took two years to develop the skill of chicken sexing. In contrast, with explicit instruction, Biederman and Shiffrar (1987) demonstrated that students could learn to do the task nearly as well in ten minutes.

⁴ There are many learning algorithms that can be applied that would produce differential learning rates. This particular example contrasting back propagation in a three layer network demonstrated a nearly two order of learning difference. This illustrates the scale of benefits that may occur in connectionist learning with and without a Control System setting up the number of vectors and the intermediate states.

So, an important characteristic of human cognition is the efficiency of learning with instruction. The ability to transfer controlled processing routines from one person to another is conceptualized in CAP2 by assuming that the architecture can instruct which stimuli to load into the modules, which produces dramatic improvements in learning rate. At present we model verbal instruction in only a very simple form. We assume that the model acquires instructions the way a human would. Through the use of language, one person can instruct another regarding the basic composition of a task, the spatial locations to be attended and monitored, the match contingencies, and the appropriate output responses. For example, in a search task, the instruction “Remember the words animal and vehicle (load lexical buffers), look at two input words above and below fixation (identify where to attend), if the words semantically match (specify that activity reports in the semantic region should be monitored), then press the ‘1’ key, else press the ‘2’ key (specify conditional responses, and the response to buffer on a match/mismatch). After a small number of rehearsals, the executive Sequential Net can similarly acquire and execute a controlled processing program reliably.

Controlled Process Monitoring of Inner Loop Messages. To execute specific goal activities the Control System needs some access to specific messages. In our original search models, controlled processing monitored only activity and priority reports on the module level (e.g., Schneider & Shiffrin, 1977). This limited the precision of control to the class of stimuli that a module coded. For example, for a module coding hunger, controlled processing could change gains and search the input space until something activated a food module. However, it could not find a specific exemplar (e.g., “food” could be coded at the module level, but not a specific exemplar like “chocolate covered almonds”)⁵. To allow controlled processing to problem solve and obtain specific goals, Schneider & Pimm-Smith (1997) proposed that **controlled processing could monitor messages on the inner loop**. This access to the inner loop provided an interpretation of the computational basis of “consciousness,” as serial messages monitored by controlled processing to achieve specific goals. This is similar to a proposal by Crick and Koch (1998) that only the top tiers of the visual system are available to consciousness.

Schneider and Pimm-Smith (1997) also proposed that **verbal instruction or “observing behavior” could allow transfer of a controlled processing routine from one person/animal to another**, greatly speeding the learning of new controlled processing procedures. In a study of how the model accommodates working memory phenomena, Schneider & Detweiler (1987) proposed a means for chunking input and output transmissions, wherein there is a parallel transmission of one module to multiple modules, and then sequential output to successive states. They also proposed the use of **fast associative learning in modules on the inner loop to enable reloading of working memory** (see McClelland, McNaughton, & O'Reilly, 1995).

Controlled processing has four very important benefits: performance of simple symbolic type tasks is possible, accurate performance can be achieved quickly in the range typically seen in human experiments (e.g., 20 trials per gate type), automaticity can emerge in the time scale seen in human learning experiments (e.g., several hundred

⁵ Even with the potential of ten thousand modules for all cognition, there would be too few module types for fine goal seeking behavior. A module could contain many vectors all mapping to a single scalar value of activity and priority. To seek specific goals, the control system needs access to some of the vector messages. In CAP2 we assume there is access to the messages on the inner loop.

trials), and learning is not limited by simple shaping and feedback, but can occur by instruction or observation.

Simulation of the transition from controlled to automatic processing. The transition from controlled to automatic processing is a result of both learning and subject strategy. Initially the subject must learn, through verbal instruction, the sequential controlled process that configures the Data Matrix to perform the task. Then the program is executed sequentially (as specified by the Control System's Sequential Net). During early practice, controlled processing generates novel combinations of module transmissions (e.g., after the visual input of "CAT" is attended and matched to a semantic category "animal", then press the "Z" key with the index finger). This results in associative learning, "cat to index finger". Priority learning causes the visual module coding "cat" to evoke a high priority output, independent of controlled processing. As automatic processing develops with consistent practice, the automatic process gradually wins out, though it remains in a horse race with controlled search. This predicts the gradual flattening of search load curves with practice (Gupta & Schneider, 1991; Schneider, 1999).

Figure 4 contrasts the slope of load curves under consistent (Figure 4B) and varied (Figure 4C) mapping conditions. Initially, automatic processing is slow and detects the target only on a subset of trials, resulting in a reduced memory or display size slope. However, as automaticity becomes reliable, the slope is eliminated. The automatic and controlled processes can operate in parallel, with the faster (automatic) process producing the response, and the subject using controlled processing as a check on the response (and perhaps to continue to strengthen the automatic processing). After several hundred trials, the subject can then drop controlled processing and maintain accuracy. Subjects can then allow automatic processing to perform the task, and can deploy controlled processing either to tune the automatic performance or to perform other tasks. However, it should be noted that subjects sometimes persist in executing unnecessary controlled processes, needlessly consuming limited controlled processing resources.⁶

CAP2 Account of Seven Phenomena of Dual Processing

Recall that the first phenomenon of dual processing is the need for extended consistent training in order to develop automatic processing. Only when there is a consistent mapping in which the target stimulus is attended to, and responded to, will the priority code increase. Recall that this occurs in the model because, when there is a sufficiently high priority code for a given input, the module can output a vector in the absence of Control system mediation. In a visual search task, after each target detection, the attended modules will have received an input, released their output vector under controlled processing, and received a global reinforcement signal. That sequence of vector-in, vector-out and reinforcement will increase the activity of the priority unit within the module and trigger associative learning (connection weight change) of the input vector to a higher priority code. For distractors, the sequence is vector-in, no vector-out, and reinforcement. In this case, the reinforcement signal without a vector-out results in an association to a lower priority code (see Figure 4A). In CM search, the target stimuli associate to ever higher priorities, and the distractors to lower priorities

⁶ In dual task training some subjects must be extensively trained to let go of controlled processing and to allocate limited this capacity controlled processing to other tasks (Schneider, & Fisk, 1982; Schumacher et al., 2001).

over repeated trials. In VM search, the priority remains at an intermediate equilibrium throughout training.

(Insert Figure 4 about here)

In CAP2, the acquisition of automaticity is slow, and needs to be slow in order to keep catastrophic interference from disrupting previously learned associations (McCloskey & Cohen, 1989). In contrast, controlled processing is set up in the SN by loading a different task vector. Since the task vector provides a differential context, catastrophic interference is prevented.

The second phenomenon of dual processing is the fast parallel performance for automatic processing and the slow serial performance for controlled processing. Automatic processing can be parallel across channels and across task stages (e.g., visual encoding, comparison, and response generation) because there are many modules that can each execute an automatic transmission without controlled gating (provided that the priority code is sufficiently high). This occurs in CM conditions (see above). Automatic processing is also parallel in memory, in that the auto-associative categorization that occurs in a module can, in parallel, map one input vector to its output vector by one settling of the module, rather than through serial comparison. In CM search, over practice, visual search comparison rates will speed up (see Figure 4B) due to automatic detections occurring before controlled process serial search identifies the target⁷.

Controlled processing in a search task is expected to be serial across items in memory, across items in the display, and across task stages. Figure 4C shows the simulated practice effects in VM search tasks. Without priority code differentiation of targets and distractors, automatic detection cannot occur. Search must stay in controlled mode even after thousands of trials of practice.

CAP2 explains the surprising inability of humans to do simple VM comparisons in parallel. With perhaps ten thousand modules comprised of tens of thousands of cells each, why should human cognition for tasks as simple as finding a letter in a visual search be so serial? In CAP2, these comparisons are made when a module receives multiple input vectors simultaneously (e.g., a memory set item vector and a probe item vector). Specifically, these input vectors are added together to obtain the values of elements in the module's input layer. The sum of these input layer elements provides a measure of the degree of correlation between the module's multiple vector inputs (see Schneider & Detweiler, 1987, Shedden & Schneider, 1990). Recall that this sum is the basis for the module's activity report. Consequently, the activity report serves as a metric of correlation or "match" between module inputs. Simulation results (Shedden & Schneider, 1990) illustrated that the accuracy of identifying a match between any two input vectors is influenced by how many total vectors arrive to the module simultaneously. The measured detection accuracy (d') was 3.6 (1% error) for serial comparisons of one visual input vector to one memory vector, and reduced to d' of 0.98 (14% error) for comparisons of two visual input vectors to one memory vector. Making a comparison of four visual inputs in parallel to one memory vector reduced d' to 0.57 (29% error). Therefore, for comparisons based on vector addition there is a need to make the comparisons serially even though the module vectors involve parallel activation of potentially thousands of units. That is, although all simulated vector operations can be

⁷ In CAP2, the Attention Controller scans items in order of priority, and therefore, as the target gets a high priority it will be the first scanned via controlled processing.

done in parallel, the criterion of having to maintain high accuracy (typically above 90% in reaction time tasks) necessitates that the network make serial comparisons.

The third phenomenon is that automatic search requires little effort and can operate in high workload situations, whereas controlled processing requires effort. In CAP2, there is only one SN, and it is capacity limited in that it must sequentially perform recurrent cycles. In contrast, automatic processing can occur in the Data Matrix with little or no SN involvement. We assume that human reporting of “effort” is based on the amount of controlled processing activity, and not the number of automatic operations (e.g., driving is reported as less effortful than VM letter search, yet driving requires far more perceptual and motor processing).

The fourth phenomena, that automatic processing is far more robust than controlled processing, is a result of the differential nature of processing in the SN and the Data Matrix. Again, the SN is a recurrent network. As associations deteriorate (e.g., due to “neural noise” introduced by alcohol) the sequential programs become unstable (e.g., making wrong branches, skipping steps, losing goals). Errors resulting from skipping steps, not maintaining criterion, or mistiming the sequencing of output gains can ultimately cause a failure to accomplish goals. In contrast, automatic processing uses feed-forward vector associations that, for well-learned patterns, are rather robust to connection loss or neural noise. In fact, one can have the situation that deactivating controlled processing (e.g., by drinking alcohol) can make the Data Matrix operate in a fully automatic mode, causing an improvement in some consistent tasks.

The fifth phenomenon, the difference in cognitive control for automatic and controlled processes, follows from the architectural structure of the model and the limitations of controlled processing. Theoretically, the lack of control is a result of the **limitation of local information for automatic processes**. In CAP2, the local connections determine the priority of the message to be sent. Because the association matrix is local to each module, priority coding occurs independent of the activities of other modules (that are not inputting to the current module) and of the activities of the Control System (unless the Control System is inputting to this specific module, as in gain control). The within-module connections that code priority can only change slowly with practice (see above), and hence it takes considerable time to unlearn an automatic process.

Controlled processing can be rapidly altered in two ways. First, a change in the task vector may alter the current SN production program (e.g., switch the task from “rhyme” judgement to “synonym” judgement). Second, control processing can maintain the contents of modules to influence later processing (e.g., maintain the lexical buffer containing the memory item such as “cat” or “cap”). The combination of task and memory buffering alters the resulting behavior (e.g., respond yes to “hat” when doing the “rhyme” task while buffering the word “cat” or while doing the “synonym” task and maintaining “cap”). Since control over these changes can occur in a single trial, the system can potentially alter behaviors that developed over decades in a matter of seconds (e.g., get up and walk on the side of your foot). The ability to control behavior by maintaining and referencing specific information provides the computational power of variable binding, as illustrated in symbolic processing systems (e.g., Newell, 1990).

The sixth phenomenon is that learning is dependent on the amount and type of controlled processing, while there is little learning in pure automatic processing.

Empirically, it was demonstrated that learning appears to be a direct function of the number of controlled processing executions, with little impact of automatic transmissions (Fisk & Schneider, 1984). Fisk and Schneider (1984) demonstrated that subjects had high recognition for controlled processing distractors presented only once, but poor recognition of distractors presented twenty times during automatic processing. In simulation studies of the development of automaticity, we examined a variety of learning rules (Gupta & Schneider, 1991). Stable learning occurred whenever there was a vector input to a module followed by a controlled process transmission (i.e., a transmission triggered by the Control System). We also tried simulations with alternative learning rules. One method allowed associative learning to occur both after automatic transmissions and controlled transmissions (and hence would predict learning during pure automatic processing). This produced a pathological state in which, as automaticity developed, many low quality inputs produced automatic outputs resulting in destructive retroactive interference that prevented the simulation from discriminating inputs it had previously classified correctly. Thus, limiting learning to controlled processing transmissions provides better stability for learned associations.

Finally, the seventh phenomenon is that the automatic attention response is dependent on the intrinsic priority of a stimulus (regardless of whether it is a target or a distractor), and is independent of the context in which the stimulus appears. This phenomenon was implemented through priority learning rules in CAP2. Without decreasing the priority following unsuccessful transmissions, simulation models can not unlearn previous priority codes (Gupta & Schneider, 1991). With both positive and negative priority learning, the simulation shows transfer for both targets and distractors (as empirically demonstrated in Shiffrin et al., 1981). Figure 4A shows the simulation of priority learning across trials. Note that a CM target has a higher priority than a VM target. Note also that an unlearned target (trial 1 CM target) has a higher priority than a well learned CM distractor. The difference in priority codes predicts the pattern of transfer as a function of previous training of targets and distractors.

Functional anatomy of controlled & automatic processing

In this section we provide a mapping between the architectural assumptions specified in CAP2 (and shown in Figures 2 and 3) to particular neural structures. As we alluded to at the onset of this paper, such mappings are complex owing to the interactive and distributed nature of processing in the cortex. Accordingly, our mapping of theory onto anatomy is necessarily speculative and preliminary, and we anticipate that future research will be informative in this effort. CAP2 was designed to predict a wide range of behavioral phenomena using a hybrid connectionist architecture. Here we examine if the qualitatively different patterns of connection and function seen in CAP2 are evident in cortical organization.

We begin by considering three basic predictions derived from the model that can be tested through functional neuroimaging research. Let us first consider predictions regarding the Control System. Recall that the Control System plays a central role during early CM task learning, but a considerably reduced role late in CM learning. Accordingly, a first prediction is that the neural substrates of the Control System should be identifiable in a contrast of brain activity produced during performance of a consistent task early in practice to that produced later in practice. More generally, areas

implementing Control System processes are expected to be active when tasks require effortful and intentional processing, and to decrease activity as automaticity develops. A second prediction stems from the fact that the Control System provides executive resources for all tasks before they are automatized. Consequently, our second prediction is that brain regions supporting controlled processing should be active in novice performance across a wide range of tasks and materials. The third prediction we consider regards the modules in the Data Matrix. One notable feature of the CAP2 model is that automatic processing is based on the operation of the same modules that were involved during earlier controlled processing of the same task. Therefore, a third prediction is that brain regions that remain active following extensive consistent practice should also have been active during early practice on the task.

(Insert Figure 5 about here)

Figure 5 shows brain areas exhibiting patterns of activity consistent with the first two predictions, and therefore implicated in Control System activity. Specifically, the regions shown exhibited substantial reductions in activation during the learning of consistently matched paired-associates, and further, showed these reductions when learning paired-associates from distinct material sets. Brain regions fitting this reduction profile include bilateral prefrontal, medial frontal (anterior cingulate), posterior parietal, occipital-temporal, and cerebellar areas. Importantly, a meta-analysis of the neuroimaging literature shows that across studies of practice-dependent change, this same network of regions is consistently reported to show reductions with practice (Chein & Schneider, in preparation). We treat these regions as candidate substrates of the Control System and consider, below, the specific control functions that each may support. Figure 6 provides a schematic diagram of these regions.

Meta-analysis also provides evidence supporting our third prediction. Namely, across studies using a baseline control condition, areas reported as active following practice are also generally reported as active early in practice. Moreover, these regions exhibiting activity both before and after practice tend to be domain specific, in that they are not generally overlapping across studies using different tasks and material. We accordingly interpret these areas of domain specific activity as reflective of processing which takes place within data modules that are relevant to the particular task being investigated.

(Insert Figure 6 about here)

The Sequential Network of CAP2 performs sequential goal-directed control, and we suggest that it maps to the dorsolateral prefrontal cortex (DLPFC). The DLPFC is frequently cited as the locus of executive processes (e.g., Roberts, Robbins, & Weiskrantz, 1998). Furthermore, axons from many different cortical regions project ultimately to the PFC (via the intralaminar nuclei of the thalamus), particularly the dorsolateral portion, thereby providing the necessary pathway through which output signals transmitted from all over the brain can converge on this centralized processor (Schneider, 1999). The DLPFC is also implicated in guiding the sequential execution of operations during controlled processing (as discussed in Schneider & Oliver, 1991), as would be expected of a region behaving as the sequential network. The SN should be active when a novel task is encountered (because a novel task requires that a new sequencing procedure be developed), and for tasks that require variable binding (e.g., VM search). Tasks that have a high workload or require the planning and ordering of many

sequential operations should also engage this brain region. Broadly, the prefrontal cortex again fits this profile. The prefrontal cortex has been shown to be active in a wide variety of cognitive tasks (Cabezza & Nyberg, 2000), and Shallice (1996) has argued that planning is one of the primary functions of this brain region. Increases in activation with workload are also often reported in the prefrontal cortex (Cohen, Perlstein, Braver, Nystrom, & et al., 1997; D'Esposito, Detre, Alsop, & Shin, 1995), and complex cognitive tasks (e.g., Raven's progressive matrices: Prabhakaran, Smith, Desmond, Glover, & Gabrieli, 1997; Tower of Hanoi: Fincham, 2002) often elicit extensive PFC activity.

In CAP2 the Attentional Control processor provides the mechanism for selective attention. We suggest that this function maps to the posterior parietal cortex (PPC). We believe that a region mediating this process should be active when a task involves frequent attentional shifts, and should be positioned near the top of brain's organizational hierarchy so that it can exert attentional gating in data modules located deeper within the hierarchy. The parietal cortex shows activation and connection patterns consistent with such a role. Specifically, neuroimaging studies have shown this region to be active in tasks that require frequent shifting of attention (Corbetta, Kincade, & Shulman, 2002; LaBar, Gitelman, Parrish, & Mesulam, 1999). Additionally, projections from the parietal cortex to the pulvinar nucleus of the thalamus, and back to other regions of the cortex, maintain a point-to-point mapping that might allow this region to exert attentional influences on processing in specific cortical modules distributed across the brain (Baizer, Desimone, & Ungerleider, 1993). The parietal cortex has also been shown to be essential for selective attention processing through neuropsychological research. For example, patients with PPC damage exhibit a "neglect syndrome" in which salient stimuli fail to attract attention (Mesulam, 2000).

The Activity Monitor watches activity from the Data Matrix and sets decision criterion, and is suggested to map to the anterior cingulate cortex (ACC). This monitoring function should be pronounced in new tasks particularly when performance is error prone and when prepotent automatic responses must be overcome. The Activity Monitor also has the responsibility of determining whether the control being exerted by the Sequential Network is sufficient to produce successful task performance (i.e., limiting errors, inhibiting inappropriate responses). Tasks that require the subject to overcome prepotent responses, such as the Stroop Task (1935), may place particular emphasis on this function. Overcoming an automatic response requires the monitoring of module comparisons and the blocking of automatic transmissions. Carter and colleagues (Carter et al., 1998; MacDonald, Cohen, Stenger, & Carter, 2000; van Veen, Cohen, Botvinick, Stenger, & Carter, 2001) have shown that the anterior cingulate cortex contributes to this type of control monitoring. Specifically, this region of cortex responds when errors are made, and when there is a high level of conflict between task demands and habitual responses. Interconnectivity between the ACC and the PFC provide a pathway through which the ACC may influence sequential controlled processes (Cohen, Botvinick, & Carter, 2000). The anterior cingulate is particularly active early in difficult tasks (Chein & Schneider, in preparation; Raichle et al., 1994). This matches an interesting characteristic of CAP2 simulations. During the first few trials of a block, the Activity Monitor must set a criterion for what activity level is classified as a match. On misses, CAP2 reduces the criterion and on false alarms it increases the criterion. If there are too many errors, the search rate is slowed to increase sensitivity. Human performance

typically shows greater errors at the beginning of a block, suggesting that some type of calibration of the decision criterion is taking place. This suggests significant interaction between the SN and Activity Monitor during early trials.

The Episodic Store records and recalls associations between the vectors in the inner loop and is suggested to map to the medial temporal lobe (hippocampus and neighboring cortex). The Episodic Store, as implemented in CAP2, provides a mechanism that enables resumption of a task after interruption through recall/reloading of inner loop modules. This function is implemented as a buffer that records the contents of inner loop processing, and can later retrieve this information if necessary. Such bufferring is likely to be mediated by the hippocampus and its adjacent neocortical areas (see discussion in Schneider & Detwiler, 1987). The hippocampus resides at the top level of the brain's organizational hierarchy, the inner loop. It receives input from association cortices responsible for high-level processing across many different domains (e.g., audition, vision, semantics, etc.). Additionally, it has been argued that the hippocampus's capacity for rapid binding of information enables this region to encode episodic memories that may then be retrieved via connections to the parahippocampal gyrus (McClelland et al., 1995).

In CAP2, the routing of report and gain control signals occurs through the Gating & Report Relay, which is suggested to map to the thalamus. The thalamus has a very extensive connection pattern to all of the other regions implicated in the Control System, as well as to other cortical areas presumed to implement the modules of the Data Matrix. The thalamus is composed of principle sensory nuclei (e.g., the pulvinar nucleus, which connects to visual sensory areas) that have point-to-point connection patterns wherein neurons from a cortical hypercolumn project down to a very tight population of neurons in the thalamus and then back-project to the hypercolumn. This projection pattern is expected of the neurons carrying priority and activity reports, as well as the output gain control scalars, in CAP2. In addition, the thalamus has specialized nuclei that connect to other cortical structures to which we have attributed Control System functions. For example, the dorsal medial thalamus projects to prefrontal cortex (SN in CAP2). Similarly, the anterior thalamic nucleus projects to the anterior cingulate gyrus and hippocampus. Moreover, reversible lesion studies of the thalamus show the importance of the thalamus as a report monitoring center that is not essential for the processing of automatic behaviors. Specifically, deactivation of the pulvinar nucleus of the thalamus (Desimone, et al, 1990) produces an inability to identify stimuli when they are in a display with competing stimuli in the visual field, but does not interfere with processing of a single stimulus in situations not requiring selective attention. This finding is consistent with removing the ability of the Control System to control the Data Matrix such that all messages must be transmitted via automatic processing (as would happen without the Gating & Report Relay). With only one stimulus to process, behavior would appear normal, while with multiple stimuli to process, the absence of selective attention would result in either no transmissions (when there are no pre-learned priorities) or multiple interfering transmissions.

In CAP2, the Control System also transmits a reinforcement signal that triggers learning in the Data Matrix. The global distribution of this reinforcement signal is suggested to map to the actions of the locus ceruleus and the ventral tegmentum. These signals differ from the previously discussed module-to-module message projections and

report signal projections that have specific targets. The diffuse projection components differ in anatomy and dynamics of operation. In CAP2, the reinforcement signal is a general broadcast signal to the Data Matrix, directing all modules that recently had a controlled process output to change their associations so that, on future trials, the input that preceded the output is more likely to evoke the output vector in a single association. The reinforcement signal would be expected to occur after attended operations that leave memory traces. Since the signal need not address individual modules (the local module activity indicates if it recently transmitted), a single neuron could effectively communicate this signal to a broad region of cortex. Again, this pattern of connectivity and dynamics is consonant with properties of the locus ceruleus (LC) and the ventral tegmentum. Projections from the locus ceruleus and the ventral tegmental area cover large regions of cortex including multiple lobes. Furthermore, neurons in this circuit predict the likelihood of reward (e.g., when a stimulus predicts the presence of a reward, the LC activity shifts from the time of the reward stimulus to the time of the predictor stimulus; Schultz, Apicella, & Ljungberg, 1993).

The locus of automatic processing in CAP2 is in the modules of the Data Matrix themselves. There are thousands of modules that process specific classes of information. Note that controlled processing activates these areas as well. The specific neuroanatomical location of these modules depends on the nature of the information being processed (e.g., specific to the sensory or motor channel conveying the information). We would expect that most areas of cortex not subserving the control network are involved in stimulus coding, and would therefore be part of the Data Matrix where automatic processing takes place. The areas we implicate in controlled processing represent a small proportion (less than 10%) of all cortex. Anatomically, we would expect the cortex structure of the modules in the Data Matrix to be similar across many regions. The proposal of a “cortical microcircuit” (Martin, 1988; Jones, 2000) is based on a repeating pattern of layering and input-output connectivity that is very similar across cortical regions, and even similar across species (from tree shrew to human). The types of cells, layering, and modular connectivity are easily likened to a large Data Matrix of modules.

To reiterate our earlier prediction from CAP2, those areas active in the final automatic task should have been active early in task acquisition as well. Controlled processing modulates the behavior of modules in the Data Matrix, thus promoting automatic component processes in those modules. Controlled processing does not develop the skill in one part of the brain and copy it to another part for automatic execution. In our neuroimaging studies of learning, we find that the areas responsible for sensory and response components in the final task were also active early in the task, and that new areas do not generally emerge (e.g., Chein & Schneider, in preparation).

There will undoubtedly be many follow-up studies relating dual processing theory to biology. The core prediction is that these studies should reveal a Control System that modulates the behavior of a Data Matrix, resulting in the development of automatic component processes. Our initial work supports this prediction.

Relationship to Biological and Computational Models

As a hybrid of symbolic, connectionist and biological models, CAP2 relates to other models in each of these classes. Here we provide a selective review of some models, and some of their points of contact with CAP2. CAP2 is fairly unique in that it places

critical emphasis on the flow of scalar report and gain control signals allowing a limited sequential controller to control a large parallel network. Norman (1969) proposed the nodes have a pertinence which is analogous to the priority report in CAP2.

Multiple investigators have suggested that attentional processing is the result of activity in a distributed network of brain regions. For example, Mesulum (1990) emphasizes that attention emerges from the interactions between the DLFPC, anterior cingulate, and posterior parietal cortex. Recall that these three structures map to the SN, Activity Monitor, and Attention Controller, respectively, in CAP2. Posner and Petersen (1990) further dissociate the component processes of attention into an anterior attention system (AAS) and a posterior attentional system (PAS). The AAS is thought to be involved in executive control, detecting targets, and task switching operations, and to include the lateral prefrontal cortex, the anterior cingulate, the amygdala and premotor areas. These functions and structures map readily to the Sequential Network and Activity Monitor in CAP2. While we have not discussed the role of the amygdala, its connection and activation patterns are consistent with it being an additional processor on the inner loop that specializes in the coding of emotional contents in memory. The PAS is assumed to play a role in attentional selection by performing the operations of disengaging, moving, and engaging attention. This posterior system is comprised of the posterior parietal cortex, the thalamus, the superior colliculus, and some cerebellar nuclei.

Laberge (1997) has also proposed a neuroanatomical framework for attention and awareness in which attention is assumed to be mediated by a triangular circuit including prefrontal areas, the thalamus, and specific posterior cortical modules that are conceptually similar to those employed in CAP2. However, in contrast to CAP2, Laberge does not make a distinction between control/report signals and representation information. For example, Laberge assumes that transmissions to the thalamus from cortical modules carry specific information regarding the actual stimuli/representations being processed (e.g., “G” is the current stimulus). Recall that in CAP2, these transmissions do not contain such stimulus specific information, but are instead believed to represent scalar codes regarding the module’s activity and priority (e.g., the current stimulus is important). We believe that evidence of preserved search task performance with single-channel, but not multi-channel, displays following deactivation of the thalamic pulvinar nucleus (Desimone et al., 1990) is more consistent with the CAP2 interpretation. That is, while Laberge’s model predicts a general degradation of performance with pulvinar deactivation, it does not predict the selective impairment of multi-channel (and not single-channel) tasks.

While the Data Matrix in CAP2 was designed to mimic the modularity of the cortex, other biological models that more explicitly detail the nature of processing taking place in cortical hypercolumns have also been developed (e.g., Martin, 1988; Goodman et al., 2001). These models address much lower-level biological mechanisms than does CAP2 (e.g., membrane channels), but share some of the same basic operating assumptions. Examples include the autoassociative properties of input units in a hypercolumn, and the ability of hypercolumn output units to be modulated by signals that scale their activity level (i.e. chandelier cells).

CAP2 also resembles several symbolic processing models. Examples of such models include SOAR (Newel, 1990), EPIC (Meyer & Kieras; 1997), ACT-R (Anderson & Lebiere; 1998), and 4CAPS (Just, Carpenter, & Varma, S. 1999). These models

assume that the nature of human cognition can be largely characterized through condition-actions rules, or productions, of the form, “if some condition is met, then perform some action.” The relevant conditions can be very complex, and can take on variable bindings. The SN in CAP2 is a simple production execution device implemented as a connectionist recurrent network. As for the SN, symbolic processing models typically assume that the production system is a limited-capacity control resource (e.g., only a single production can be executed per cycle, regardless of the information it processes). Further, like CAP2, most production system models assume that limited information is available to the production system. Recall that in CAP2, the SN cannot operate on messages processed below the association tier of the Data Matrix. This is similar in concept to the limited input working memory in EPIC and to buffers in ACT-R 5.0 (Anderson et al., submitted). Relatedly, the inner loop modules in CAP2 map fairly readily to the peripheral processors in EPIC (auditory, visual, ocular motor, visual motor, vocal motor, manual motor, tactile) and to buffers in ACT-R 5.0 (goal, visual, manual, imaginal, retrieval). Likewise, the Episodic Store in CAP2 relates to the declarative procedure store in EPIC and the retrieval buffer in ACT-R. Work on ACT-R 5.0 (Anderson et al., submitted) also includes predictions regarding the neuroanatomical locations of production system components (e.g., goal buffer to DLPFC, retrieval buffer to VLPFC, productions to the basal ganglia) that resemble the mappings that we have proposed between CAP2 and the brain.

Despite these similarities, there are important conceptual differences between production models and CAP2. Production systems typically learn by building new productions, strengthening old ones, and converting productions from declarative form to procedural form. In procedural form, productions are faster and can execute in parallel (e.g., Schumacher et al., 2001). However, precisely when and how such learning can take place is typically not addressed. In CAP2, instead of productions becoming compiled and still executed by the control system, modules in the data matrix become automatic and operate in the absence of control system input. This allows parallel processing as long as the inputs are clear, responses are consistent, and the executive releases sequencing to the modules.

The CAP2 architecture utilizes multiple connectionist concepts. The basic multilayer processing typical of PDP models (McClelland & Rumelhart, 1988) is present in each module. The fast learning of the Episodic Store uses techniques developed in models exploring episodic functions of the hippocampus (McClelland, McNaughton & O'Reilly 1995). The sequential recurrent nets developed by Elman (1990) were the basis for the SN in the Control System. However, CAP differs in that it assumes part of the brain has evolved to control information flow in other parts of the brain. CAP2 also assumes modules have a wider variation of unit types than is typical in connectionist layers (e.g., report cells, gain control units).

In summary CAP2 was developed to predict the the dual process nature of human processing and skill acquisition utilizing biologically reasonable component architectures. It provides a rich set of behavioral, computational, and biological predictions that can be related to other models and the developing physiological evidence. As modeling becomes more constrained by the growing behavioral and imaging data, we see a convergence of operations and cortical structures across modeling methods.

Conclusion

Dual processing theory has seen very productive development in the last twenty-five years. It provides quantitative predictions of the dramatic changes in performance that emerge with task practice and a suitable interpretation of the cortical changes that coincide with learning. It predicts that in consistent, but not varied, tasks, automatic component processes will develop. A wide range of behavioral results are coherently interpreted within this framework. Automaticity leads to fast, parallel, robust, low effort performance, but requires extended training, is difficult to control, and shows little memory modification. In contrast, controlled processing is slow, serial, effortful, and brittle, but it allows rule-based processing to be rapidly acquired, can deal with variable bindings, can rapidly alter processing, can partially counter automatic processes, and speeds the development of automatic processing. The benefits of dual processing were described initially by James in 1890 and detailed extensively by Schneider and Shiffrin in 1977, almost a century later. James (1890) remarked that, “the more of the details of our daily life we can hand over to the effortless custody of automatism, the more our higher powers of mind will be set free for their own proper work.” In our work, we have not only shown how transitions between processing modes may occur, but have extended James’s view by recognizing that controlled and automatic processing are complementary modes of behavior supported by different processing architectures.

There is a developing synergism between the behavioral, computational, and biological interpretations of dual processing theory. We have described a computational model of human performance in the controlled and automatic process modes, which also explains the transitions between the modes. The modeling work has lead to an understanding of the relevant processing tradeoffs, and to elaboration of the expected computational architecture. A rapidly advancing brain imaging literature provides a forum for tests of this model, and offers initial support for key concepts in dual processing theory. The physiology has also stimulated the implementation of new concepts in the model, such as the presence of inner-loop communications and declarative memory coding of inner loop information. The early work on dual processing theory has evolved dramatically over the past twenty five years, and is expected to help provide interpretations for phenomena revealed in the developing biological and computational fields for many years to come.

References

- Anderson, J.A. (1983). Cognitive and psychological computation with neural models. IEEE, Transactions on systems, man, and cybernetics, SMC-13, 799-815.
- Anderson, J. R. (1992). Automaticity and the ACT* theory. Am J Psychol, 105(2), 165-180.
- Anderson, J. R. & Lebiere, C. (1998). The atomic components of thought. Mahwah, NJ: Erlbaum.
- Anderson, J. R., Bothell, D., Byrne M. D. & Lebiere, C. (submitted). An Integrated Theory of the Mind.
- Atkinson, R.C. & Shiffrin, R.M. (1968). Human memory: A proposed system and its control processes. In K.W. Spence & J.T. Spence (Eds.), The psychology of learning and motivation: Advances in research and theory, 2. New York: Academic Press.
- Baizer, J. S., Desimone, R., & Ungerleider, L. G. (1993). Comparison of subcortical connections of inferior temporal and posterior parietal cortex in monkeys. Vis Neurosci, 10(1), 59-72.
- Bargh, J. A. (1992). The ecology of automaticity: toward establishing the conditions needed to produce automatic processing effects. Am J Psychol, 105(2), 181-199.
- Biederman, I., & Shiffrar, M. (1987). Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual-learning task. J Exp Psychol Learn Mem Cogn, 13(4), 640-645.
- Briggs, G. E., & Johnsen, A. M., (1973) On the natyure of central processsing in choice reactions. Memory and Cognition. 1, 91-100
- Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. J Cogn Neurosci, 12(1), 1-47.
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. Science, 280(5364), 747-749.
- Chein, J. M., & Schneider, W. (in preparation). Evidence of a domain general learning network: An FMRI investigation with verbal and nonverbal paired-associates.
- Cleeremans, A., Servan-Schreiber, D., & McClelland, J.L. (1989.) Encoding sequential structure in simple recurrent networks. Technical report CMU-CS-88-183, Computer Science Department. Carnegie Mellon University, Pittsburgh, PA.
- Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control and schizophrenia: recent developments and current challenges. Philos Trans R Soc Lond B Biol Sci, 351(1346), 1515-1527.
- Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., & et al. (1997). Temporal dynamics of brain activation during a working memory task. Nature, 386(6625), 604-608.
- Corbetta, M., Kincade, J. M., & Shulman, G. L. (2002). Neural systems for visual orienting and their relationships to spatial working memory. J Cogn Neurosci, 14(3), 508-523.
- Crick F., Koch, C. (1998) Consciousness and Neuroscience Cerebral Cortex, 8:97-107
- Desimone, R., Wessinger, M., Thomas, L. & Schneider, W. (1990). Attentional control of visual perception: Cortical and subcortical mechanisms. In Cold Spring

Harbor Symposium on Quantitative Biology, Vol. LV. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press.

D'Esposito, M., Detre, J. A., Alsop, D. C., & Shin, R. K. (1995). The neural basis of the central executive system of working memory. Nature, 378(6554), 279-281.

Elman, J.L. (1990). Finding structure in time. Cognitive Science, 14:179-212.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. Cereb Cortex, 1(1), 1-47.

Fincham, J. M., Carter, C. S., van Veen, V., Stenger, V. A., & Anderson, J. R. (2002). Neural mechanisms of planning: a computational analysis using event-related fMRI. Proc Natl Acad Sci U S A, 99(5), 3346-3351.

Fisk, A. D. & Schneider, W. (1981). Control and automatic processing during tasks requiring sustained attention: A new approach to vigilance. Human Factors, 23(6), 737-750.

Fisk, A. D. & Schneider, W. (1982). Type of task practice and time-sharing activities predict performance deficits due to alcohol ingestion. Proceedings of the Human Factors Society, 926-930.

Fisk, A. D., & Schneider, W. (1983). Category and word search: generalizing search principles to complex processing. J Exp Psychol Learn Mem Cogn, 9(2), 177-195.

Fisk, A. D., & Schneider, W. (1984). Memory as a function of attention, level of processing, and automatization. J Exp Psychol Learn Mem Cogn, 10(2), 181-197.

Goodman, P. H., Courtenay Wilson E., Maciokas, J. B., Harris, F. C., Gupta, A. G., Louis, J. L., & Markram, H. (2001) Large-Scale Parallel Simulation of Physiologically Realistic Multicolumn Sensory Cortex NIPS

Gupta, P. & Schneider, W. (1991). Attention, automaticity, and priority learning. In Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society (pp.534-539). Hillsdale, NJ: Erlbaum.

Hancock, P. A. (1986). The effect of skill on performance under an environmental stressor. Aviat Space Environ Med, 57(1), 59-64.

Heuer, H., Spijkers, W., Kiesswetter, E., & Schmidtke, V. (1998). Effects of sleep loss, time of day, and extended mental work on implicit and explicit learning of sequences. J Exp Psychol Appl, 4(2), 139-162.

James, W. (1890). The principles of Psychology (Vol. 1). New York: Holt.

Johnsen, A. M., & Briggs, G. E. (1973). On the locus of display load effects in choice reactions. J Exp Psychol, 99(2), 266-271.

Jones, E. G. (2000). Microcolumns in the cerebral cortex. Proc Natl Acad Sci USA; 97, 5019-21.

Just, M. A., Carpenter, P. A., & Varma, S. (1999). Computational modeling of high-level cognition and brain function. Hum Brain Mapp, 8(2-3), 128-136.

Kristofferson, M. (1972). Effects on practice on character classification performance. Canadian Journal of Psychology 26: 54-60.

LaBar, K. S., Gitelman, D. R., Parrish, T. B., & Mesulam, M. (1999). Neuroanatomic overlap of working memory and spatial attention networks: a functional MRI comparison within subjects. Neuroimage, 10(6), 695-704.

LaBerge, D. (1997). Attention, awareness, and the triangular circuit. Consciousness and Cognition, 6, 149-181.

Lebiere, C. & Anderson, J. R. (1993). A connectionist implementation of the ACT-R production system. In Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society, pp. 635-640.

Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: theory and data. Cognit Psychol, 12(4), 523-553.

Logan, G. D. (1992). Attention and preattention in theories of automaticity. Am J Psychol, 105(2), 317-339.

MacDonald, A. W., 3rd, Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. Science, 288(5472), 1835-1838.

Martin, K. A. (1988). The Wellcome Prize lecture. From single cells to simple circuits in the cerebral cortex. Q J Exp Physiol, 73(5), 637-702.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev, 102(3), 419-457.

McClelland, J.L. & Rumelhart, D.E. (1988). Explorations in parallel distributed processing: A handbook of models, programs, and exercises. Cambridge, MA: MIT Press.

McCloskey, M. & Cohen, N.J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (Ed.) The psychology of learning and motivation (Vol. 24, pp.109-166).

Mesulam, M. M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language, and memory. Annals of Neurology, 28, (5), 597-613.

Mesulam, M. M. (2000). Attentional networks, confusional states, and neglect syndromes. In Principles of behavioral and cognitive neurology (2nd ed., pp. 174-256). London: Oxford University Press.

Meyer, D. E. & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance. Part I. Basic mechanisms. Psychological Review, 104, 2-65.

Miller, E. K. (2000). The prefrontal cortex and cognitive control. Nature Reviews Neuroscience, 1, 59-65.

Miller, E. K. Cohen J. D. (2001) an integrative theory of prefrontal cortex function Annual Review of Neuroscience. 24:167–202

Miyake A. & Shah B. (Eds.) (1999) Models of working memory: Mechanisms of active maintenance and executive control, pp 340-374. Cambridge, UK: Cambridge University Press.

Naatanen, R. (1992). Attention and brain function. Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

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

Norman (1969) Memory and Attention New York, John Wiley & Sons

O' Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically based computational model of working memory, Miyake, Akira (Ed); Shah, Priti (Ed). (1999). Models of working memory: Mechanisms of active maintenance and executive control. (pp. 375 411). New York, NY, US: Cambridge University Press.

Pashler, H., Johnston, J. C., & Ruthruff, E. (2001). Attention and performance. Annu Rev Psychol, 52, 629-651.

Posner, M. I. Petersen, SE (1990) The attention system of the human brain. Annual Review of Neuroscience 13: 25-42

Prabhakaran, V., Smith, J. A., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. (1997). Neural substrates of fluid reasoning: an fMRI study of neocortical activation during performance of the Raven's Progressive Matrices Test. Cognit Psychol, 33(1), 43-63.

Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A. M., Pardo, J. V., Fox, P. T., & Petersen, S. E. (1994). Practice-related changes in human brain functional anatomy during nonmotor learning. Cereb Cortex, 4(1), 8-26.

Roberts, A.C., Robbins, T.W., & Weiskrantz, L. (1998). The prefrontal cortex: Executive and cognitive functions. Oxford: Oxford University Press.

Schneider, W. (1999). Working memory in a multilevel hybrid connectionist control architecture (CAP2), Miyake, Akira (Ed); Shah, Priti (Ed). (1999). Models of working memory: Mechanisms of active maintenance and executive control. (pp. 340 374). New York, NY, US: Cambridge University Press.

Schneider, W., & Detweiler, M. (1987). A connectionist/control architecture for working memory, Bower, Gordon H. (Ed). (1987). The psychology of learning and motivation: Advances in research and theory, Vol. 21. (pp. 53 119). San Diego, CA, US: Academic Press, Inc.

Schneider, W., & Fisk, A. D. (1982). Degree of consistent training: improvements in search performance and automatic process development. Percept Psychophys, 31(2), 160-168.

Schneider, W. & Oliver, W. L. (1991). An instructable connectionist/ control architecture: Using rule-based instructions to accomplish connectionist learning in a human time scale. In K. Van Lehn (Ed.), Architectures for intelligence: The 22nd Carnegie Mellon symposium on cognition (pp.113-145). Hillsdale, NJ: Erlbaum.

Schneider, W. & Pimm-Smith, M. (1997). Consciousness as a message aware control mechanism to modulate cognitive processing. Chapter in J. Cohen & J. Schooler (Eds.) Scientific Approaches to Consciousness: 25th Carnegie Symposium on Cognition, Erlbaum Assoc., Mahwah, NJ. 65-80.

Schneider, W., Pimm-Smith, M., & Worden, M. (1994). Neurobiology of attention and automaticity. Curr Opin Neurobiol, 4(2), 177-182.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. Psychological Review, 84(1), 1-66.

Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. J Neurosci, 13(3), 900-913.

Schumacher, E. H., Seymour, T. L., Glass, J. M., Fencsik, D. E., Lauber, E. J., Kieras, D. E., & Meyer, D. E. (2001). Virtually perfect time sharing in dual-task performance: Uncorking the central cognitive bottleneck. Psychological Science, 12, (2), 101-108.

Shallice, T., & Burgess, P. (1996). The domain of supervisory processes and temporal organization of behaviour. Philos Trans R Soc Lond B Biol Sci, 351(1346), 1405-1411; discussion 1411-1402.

Shedden, J. M. & Schneider, W. (1990). A connectionist model of attentional enhancement and signal buffering. in Proceedings of the Twelfth Annual Conference of the Cognitive Science Society (pp. 566-573). Hillsdale, NJ: Erlbaum.

Shiffrin, R. M. (1988). Attention. In R. C. Atkinson & R. J. Herrnstein & G. Lindzey & R. D. Luce (Eds.), Steven's handbook of experimental psychology, vol 2: learning and cognition. (pp. 739-811). New York, NY: John Wiley & Sons.

Shiffrin, R. M., Dumais, S. T. & Schneider, W. (1981). Characteristics of automatism. In J. Long & A. Baddeley (Eds.), Attention and performance IX (pp. 223-238). Hillsdale, NJ: Erlbaum.

Shiffrin, R. M., & Gardner, G. T. (1972). Visual processing capacity and attentional control. J Exp Psychol, 93(1), 72-82.

Shiffrin, R. M., McKay, D. P., & Shaffer, W. O. (1976). Attending to forty-nine spatial positions at once. J Exp Psychol Hum Percept Perform, 2(1), 14-22.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. Psychological Review, 84(2), 127-190.

Stanovich, K. E. (1987). The impact of automaticity theory. J Learn Disabil, 20(3), 167-168.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. J Exp Psychol, 18, 643-662.

Townsend, J.T. & Ashby, F.G. (1983). Stochastic modeling of elementary psychological processes. New York, NY: Cambridge University Press.

Touretzky, D. S. (2002) Connectionist and symbolic representations. In M. A. Arbib (ed.), Handbook of Brain Theory and Neural Networks, 2nd edition. MIT Press.

Treisman, A., & Souther, J. (1985). Search asymmetry: a diagnostic for preattentive processing of separable features. J Exp Psychol Gen, 114(3), 285-310.

Van Veen, V., Cohen, J. D., Botvinick, M. M., Stenger, V. A., & Carter, C. S. (2001). Anterior cingulate cortex, conflict monitoring, and levels of processing. Neuroimage, 14(6), 1302-1308.

White, E. L (1989). Cortical Circuits Synaptic Organization of the cerebral cortex. Structure, function and Theory Boston Birkhauser.

Wolfe, J.M. (2001). Guided Search 4.0: A guided search model that does not require memory for rejected distractors [Abstract]. Journal of Vision, 1(3), 349a, <http://journalofvision.org/1/3/349>, DOI 10.1167/1.3.349

Worden, & Schneider, (1995) Cognitive task design for fMRI International Journal of Imaging Science & Technology, 6, 253-270.

Young, R. M., & Lewis, R. L. (1999). The Soar cognitive architecture and human working memory, Miyake, Akira (Ed); Shah, Priti (Ed). (1999). Models of working memory: Mechanisms of active maintenance and executive control. (pp. 224 256). New York, NY, US: Cambridge University Press.

Table 1

Summary of CAP2 Components and Features

- **Data Matrix of Modules**
 - **Module Internal structure and operations**
 - Two layer connectionist networks with autoassociative input layer and gain modulated output layer
 - Learning types: *Input-Output Associative* and *Priority Coding*
 - **Module Inputs**
 - Vector input from other modules/systems
 - Scalar control inputs: *Output Gain*, *Feedback Gain* and *Reinforcement*
- Signal*
 - **Module Outputs**
 - Vector output to other modules/systems
 - Scalar outputs: *Activity Report* and *Priority Report*
 - **Data Matrix Organization**
 - Modules are organized into tiers, and tiers into regions
 - Output vector transmissions occur bi-directionally from tier to tier
 - Modules in the top level tiers of each region connect via an *Inner Loop* and send vector outputs to the Control System
 - *Peripheral Tier* modules receive vector inputs from sensory receptor systems or send vector outputs to motor systems
- **Control System**
 - **Comprised of five processors**
 - *Sequential Net* performs sequential control programs
 - *Attention Controller* tracks module priority signals and changes output gains
 - *Activity Monitor* monitors module activity reports
 - *Gating & Report Relay* routes report and gain control signals
 - *Episodic Store* records associations between the vectors in the inner loop
 - **Inputs to Control System**
 - Module Priority and Activity Reports
 - Inner Loop module vector outputs
 - **Outputs from Control System**
 - Output Gain control to a small number of modules
 - Feedback Gain control to a particular region
 - Global Reinforcement Signal that triggers associative and priority learning
 - Vector messages to the Inner Loop
- **Processing modes**
 - *Automatic Process* – when a module receives a message with a high priority code and transmits the associated output vector in the absence of a Control System input

- **Controlled Process** – when a module must wait for an output gain signal from the Control System in order to transmit its output vector

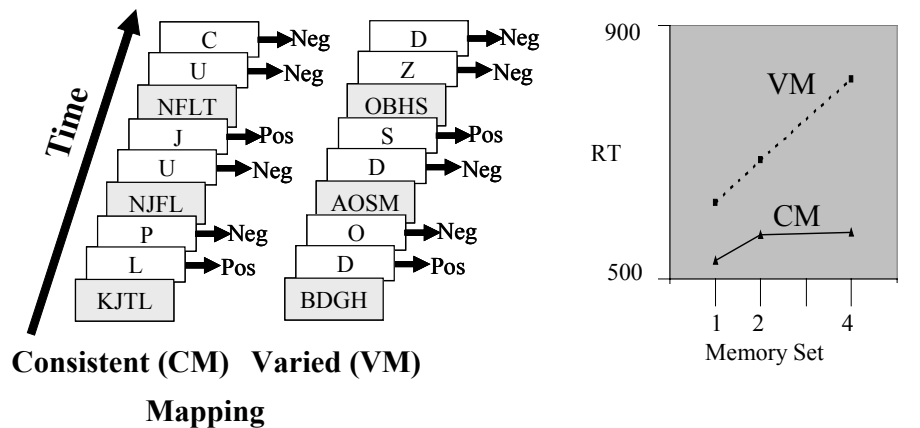


Figure 1. A visual search task. The memory sets for each trial are shown in shaded frames. The subject memorizes the list. Then, single probe letters are presented sequentially. The subject presses the positive key if the probe matches any of the memory set letters. If the probe does not match, the negative key is pressed. In CM, a letter appearing in the memory set (a target) is never shown as a distractor on subsequent trials. In VM, the same letter can be a target on one trial, and distractor on the next. The figure insert shows typical reaction time results (adapted from Briggs & Johnsen, 1973).

Module & Regional Structure

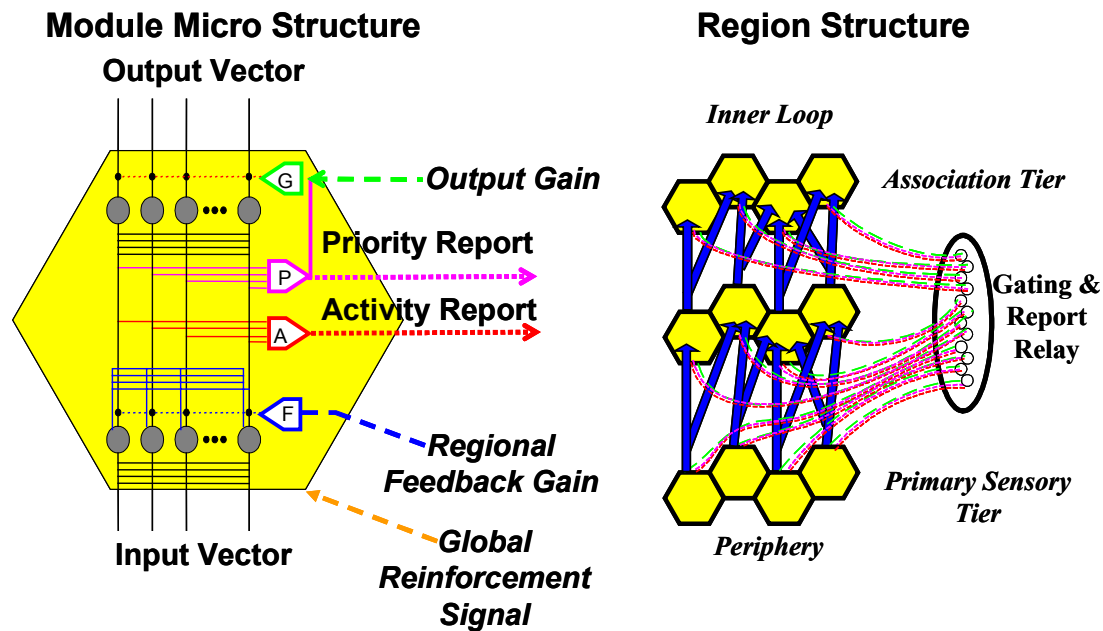


Figure 2. (LEFT) Microstructure of a CAP2 module. A set of input vectors arrive to the module (bottom left). These vectors pass through a connection matrix that reflects prior learning, and evoke a new activation pattern in the module's input layer. The input layer then activates the output layer through a second connection matrix. An auto-associative loop connects the input layer units back to other units of the input layer, which provides feedback gain modulated by the feedback gain unit (F). The module output is controlled by an output gain unit (G), which modifies the ability to send a transmission to other modules. There are two output report units, the Activity Report unit (A) and the Priority Report unit (P), that are communicated to the Control System. The Priority Report also has a local impact on the output gain unit, which allows the module to transmit, even in the absence of a Control System input, when a high priority signal is present. There is also a global reinforcement signal that is broadcast to all modules, which causes recently transmitting modules to change their connection weights. (RIGHT) The organization of modules in a single region. These modules are layered into tiers, which represent the hierarchical arrangement of processing from the periphery (primary sensory tier) to the inner loop (association tier). Each module projects to multiple modules in the next tier, and communicates its report and output gain signals via the Gating & Report Relay.

Macro Level Structure

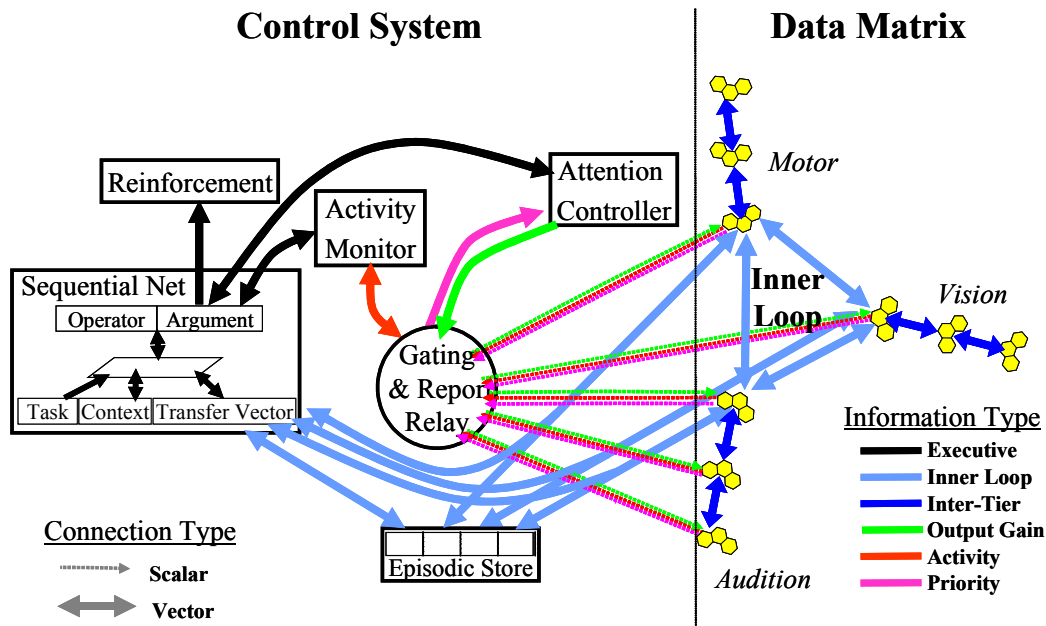


Figure 3. Macro-Structure of the CAP2 model. The Control System is shown on the left and the Data Matrix on the right. The Data Matrix is comprised of multiple modules, organized into tiers and regions, sending vector messages to the next level within their region, and scalar report signals to the Gating & Report Relay. The Inner Loop includes the top tier modules of each region (association areas). These connect to other association regions, to Episodic Store, and to the Sequential Net. . The Sequential Net (left) is a sequential recurrent network executing serial control programs. The Sequential Net outputs commands to the Activity Monitor, Attention Controller, and Episodic Store processors.

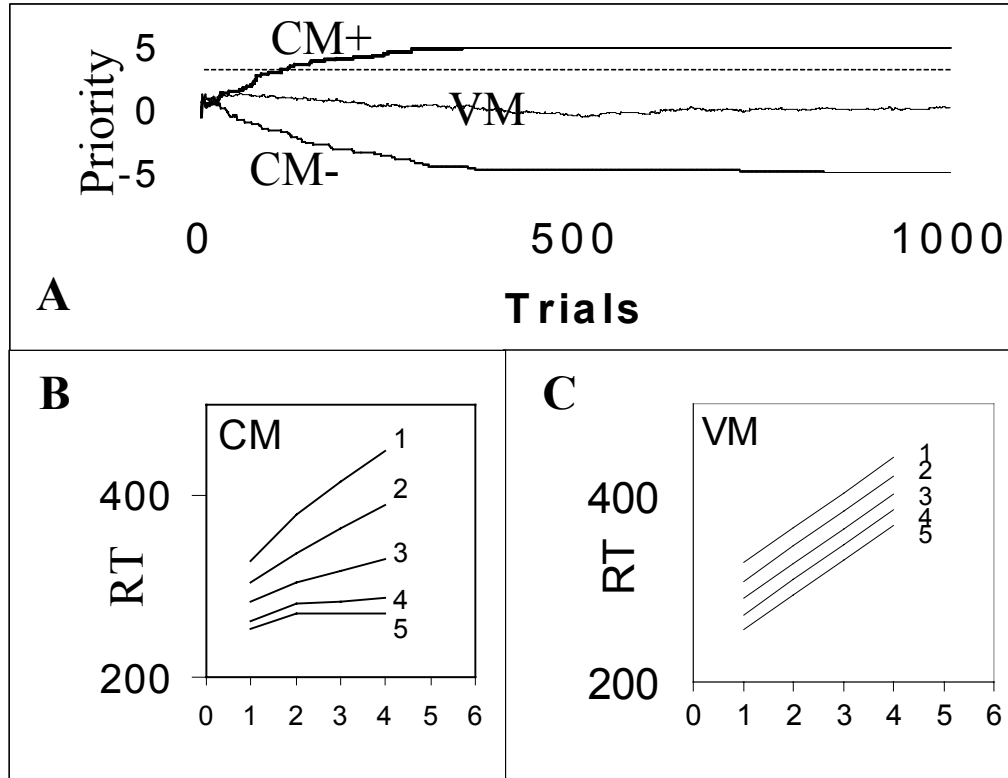


Figure 4. Priority learning simulation. (A) Priority level across repeated trials for CM target stimuli (CM+), CM distractor stimuli (CM-), and VM stimuli (targets and distractors). The dashed line represents the automatic transmission threshold. (B) Simulated CM search results for number of comparisons. Each line represents the average performance of the model after each increment of 50 trials. (C) Simulated VM search results. (Adapted from Gupta & Schnieder, 1991).

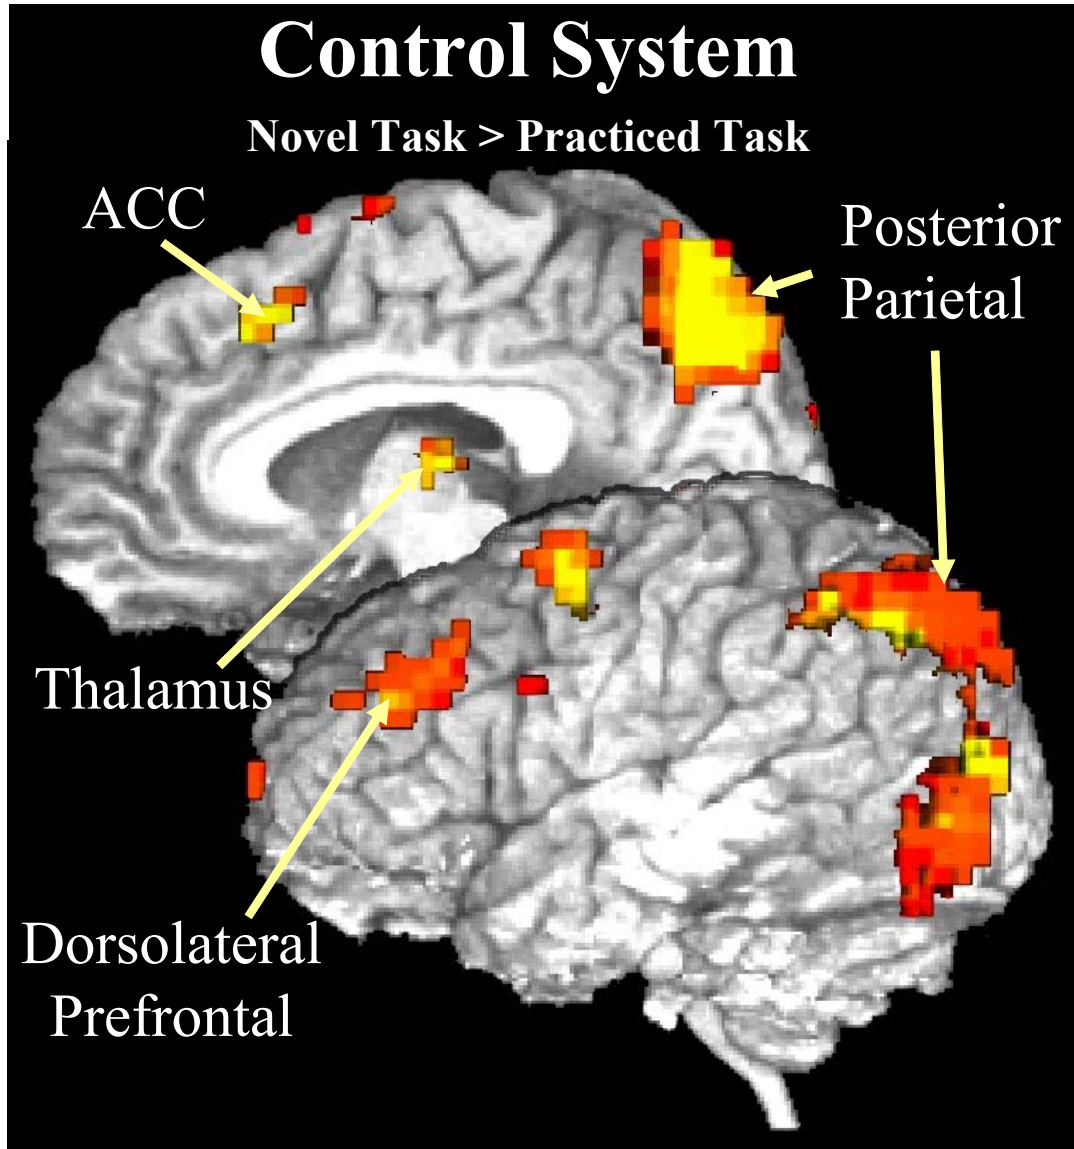


Figure 5. fMRI image of brain regions exhibiting decreased activity with practice in a consistent visual paired-associate learning task. The colored areas indicate significant areas where activation was decreased in practiced relative to the novel performance.

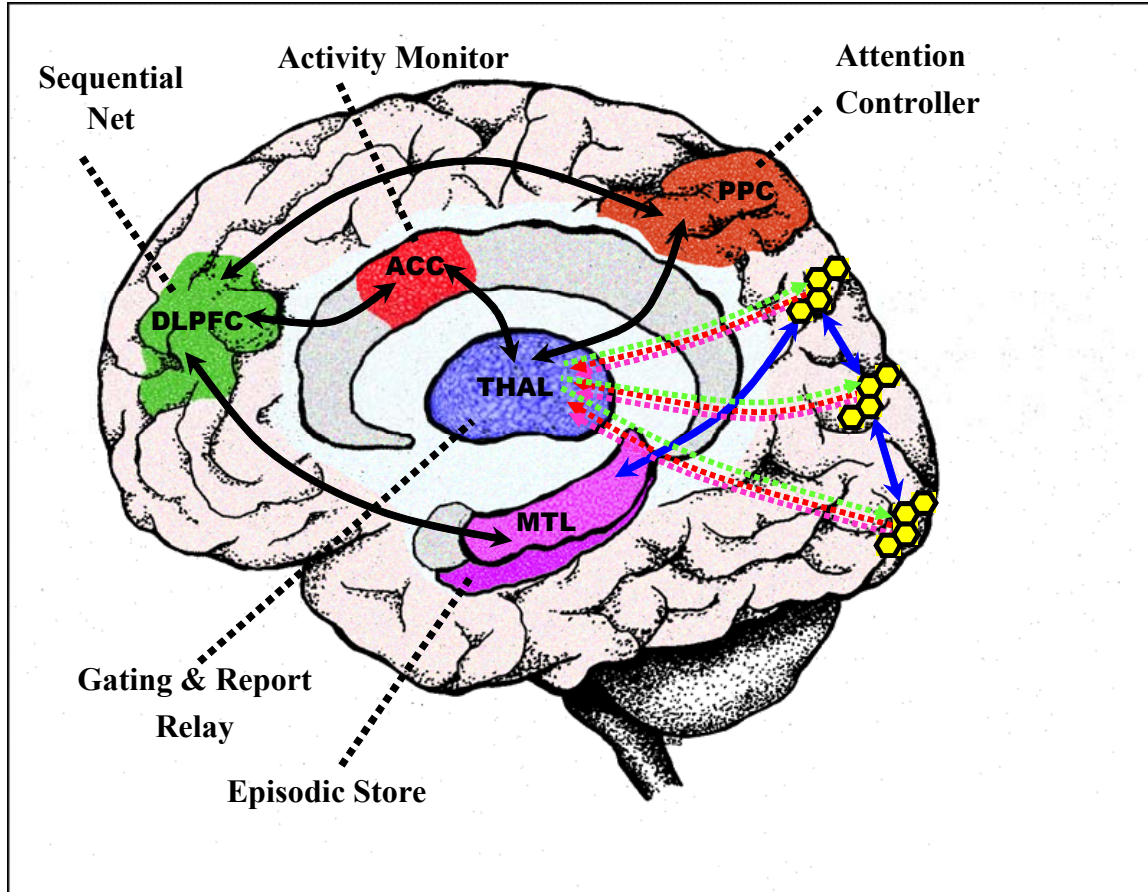


Figure 6. Mapping of CAP2 macro scale architecture to brain areas. The executive Sequential Net is assumed to be located in dorsolateral prefrontal cortex (DLPFC). The Attention Controller maps to posterior parietal cortex (PPC) and the Activity Monitor to anterior cingulate cortex (ACC). The Episodic Store maps to the medial temporal lobe (MTL), including the structures of the hippocampal complex. The Gating & Report Relay maps to the thalamus (THAL), with different thalamic nuclei connecting to alternative Control System processors, receiving report signals from the data matrix modules, and sending output gain signals to the modules. The arrows between regions illustrate known anatomical pathways. Shown on the right of the figure are sample modules in the visual region of the Data Matrix, with report and control signals form each tier connecting to the thalamus.