

A Physiologically Based System Theory of Consciousness

by *L. Andrew Coward*

email landrewc@aol.com

Abstract

A system which uses large numbers of devices to perform a complex functionality is forced to adopt a simple functional architecture by the needs to construct copies of, repair, and modify the system. A simple functional architecture means that functionality is partitioned into relatively equal sized components on many levels of detail down to device level, a mapping exists between the different levels, and exchange of information between components is minimized. In the instruction architecture functionality is partitioned on every level into instructions, which exchange unambiguous system information and therefore output system commands. The von Neumann architecture is a special case of the instruction architecture in which instructions are coded as unambiguous system information. In the recommendation (or pattern extraction) architecture functionality is partitioned on every level into repetition elements, which can freely exchange ambiguous information and therefore output only system action recommendations which must compete for control of system behavior. Partitioning is optimized to the best tradeoff between even partitioning and minimum cost of distributing data. Natural pressures deriving from the need to construct copies under DNA control, recover from errors, failures and damage, and add new functionality derived from random mutations has resulted in biological brains being constrained to adopt the recommendation architecture. The resultant hierarchy of functional separations can be the basis for understanding psychological phenomena in terms of physiology. A theory of consciousness is described based on the recommendation architecture model for biological brains. Consciousness is defined at a high level in terms of sensory independent image sequences including self images with the role of extending the search of records of individual experience for behavioral guidance in complex social situations. Functional components of this definition of consciousness are developed, and it is demonstrated that these components can be translated through subcomponents to descriptions in terms of known and postulated physiological mechanisms.

1. Introduction

A scientific theory of consciousness must have a number of characteristics. It must propose a one to one correspondence between psychological and physiological states. For example any psychological state X must have a corresponding physiological state x . Causal connections between physiological states at the physiological level of description must exist whenever there are causal connections between psychological states at the psychological level of description. For example, if physiological states x and y correspond with psychological states X and Y and state X causes state Y at the psychological level of description, then state x must cause state y at the physiological level of description. The same physiological state x cannot correspond with multiple psychological states $X1, X2, X3$, etc. although multiple physiological states could correspond with the same psychological state. For example, experiences of the colors red and blue cannot correspond with the same physiological states. Differences between experiences at the psychological level between, for example, different individuals, must correspond with differences at the physiological level in a consistent fashion, i.e. when there is a difference between the experience of the color red described at the psychological level and these differences are observed in a number of instances, the corresponding physiological states must differ in a consistent fashion. If the difference in the feel of the color red between two individuals can be described by an outside observer then it must be reflected in a difference at the physiological level. An individually unique feel must therefore correlate consistently with an individually unique physiological state.

Electronic systems have been designed and built which can perform extremely complex combinations of functions using large numbers of components. Such systems employ thousands of millions of transistors to

perform combinations of thousands of interacting features. The design process for such systems involves creation of a functional architecture in which functionality at high level is partitioned into components, and components partitioned into more and more detailed subcomponents through a series of levels down to the device level. At each level the same functionality is described, but at a different degree of detail. Such a functional architecture, if it existed in biological brains, would therefore be the basis for a scientific theory of consciousness as described in the previous paragraph. Design experience with electronic systems has demonstrated that unless a functional architecture exists and is simple in the sense that components on one level are roughly equal in size and require limited information exchange to perform their functions, the resulting systems are extremely difficult to build, repair, or modify. Coward (1990, 1997a) has argued that similar requirements result in natural selection exerting very strong pressures towards simple functional architectures in biological brains.

Commercial systems are designed by partitioning functionality into components which are conceptually instructions that generate commands to the system. To be able to generate system commands, all the information input to components must be meaningful and unambiguous. A critical distinction in system design is between global and local data. Global data is meaningful and unambiguous throughout the system. Local data is only unambiguous within one functional component and its subcomponents. The terms local and global define logical data types, both types could have wide physical distribution within a system. To achieve a simple functional architecture, the design process for commercial systems optimizes the functional partitioning to minimize the need to convert local data to global data.

The unambiguous value of an element of global data could be required as input by any functional component. There is therefore a need for a reference copy of global data, which is generally stored in a subsystem called memory. Data which is unambiguous within a component and its subcomponents but ambiguous outside the component will also be accessed by different subcomponents via memory. If one component is operating on elements of global information, those elements are not unambiguously defined for other components. Components can therefore only access global information sequentially, and the currently active component is generally executed in a subsystem called processing. The memory, processing separation is therefore ubiquitous in commercial systems, and sequential execution can only be avoided if global data can be partitioned into orthogonal segments, each of which is adequate for a different functional component. The von Neumann architecture is an important instance of the instruction based functional architecture in which many of the instructions are regarded as unambiguous global data and recorded in memory.

The process of functional separation into components which exchange unambiguous information is a paradigm which is deeply embedded in electronic design. Even design using neural networks aims to produce modules which generate unambiguous outputs such as specific vectors for well defined cognitive categories. Dependence on exchange of unambiguous information between functional modules results either in memory, processing separation and sequential execution or in an inability to construct a system to perform a complex functionality. Coward (1990) proposed that the only way to avoid this constraint is to allow free exchange of ambiguous information between functional components. Such components are programmed to detect specific combinations of input information, but because the information is ambiguous, although not meaningless, the outputs of such components can only be system action recommendations rather than commands, and the recommendations must compete for control of system behavior. The specific combinations are not patterns in the full cognitive sense. The combinations of data are ambiguous, and therefore even if an approximately equivalent cognitive condition could be found, would not always repeat when the cognitive condition is present, and would sometimes repeat when it is not present. Such repetitions are ambiguous indications that a particular system response is appropriate. Consistent application of this paradigm results in a system with the recommendation architecture in which the major subsystem separation is into repetition similarity clustering and competition rather than the memory and processing found in instruction architectures.

Coward (1990) argued that the recommendation architecture was an important alternative to the instruction architecture for systems using large numbers of devices to perform a complex functionality, and that biological brains exhibit a striking resemblance to systems with the recommendation architecture. Coward (1998) proves that these two architectures are the only options for such systems, the argument is summarized in the appendix to this paper. The understanding of consciousness in terms of physiology requires a consistent functional architectural approach to the conscious system as a whole. Such an approach must be based on an understanding of the architectural constraints which apply to any system performing very complex combinations of functionality.

2. Architectural Constraints On Complex Systems

The functionality of a system is the behavior which it generates using information derived both internally and from its environment. For example, an electronic system might take inputs from a keyboard, internal memory, and data communication links and perform a range of computing and display tasks such as word and graphics processing and email. The problem in system design is to take this high level definition of the system functionality and create a description of a system which will perform as defined in terms of how its constituent devices (e.g. transistors) are connected together and organized. For many systems the functionality is extremely complex in the sense that many different but interacting functions are performed, and the number of devices required is very large. For example, a single telecommunications central office switch may require billions of transistors to provide the many interacting functions required to deliver reliable telephone service to 100 thousand users.

Any system which uses large numbers of devices to perform a complex functionality is subject to severe architectural constraints. To understand the nature of these constraints, imagine a design process in which (in a caricature of the process of biological evolution) a large number of technicians were given a large number of devices and told to connect at random. Periodically the result would be tested, and eventually a system found which performed as required. There are some severe problems with such a system. There are no blueprints which can guide the process of building another copy: the only option is to duplicate the original, device by device, connection by connection. If an error were made in such a duplication process a functional problem would result, but there would be no easy way to use the knowledge of the functional problem to identify and correct the error. Similarly, if a device failed during operation it would be very hard to identify which device was defective. Finally, device level changes would generate complex and unpredictable functional changes, and if there were a need to modify the functionality in a controlled fashion there would be no way to identify what device level changes would produce the desired functional modifications.

Such a system is therefore forced to adopt a simple functional architecture by these needs to build copies of the system, to recover from construction errors and component failures, to configure for individual system conditions, to economize on resources through sharing of subsystems across multiple functions, and to add new features (Coward 1990, 1997a). In a system with a functional architecture, high level functionality is partitioned into components, components into subcomponents, and so on down to the functionality of the individual devices from which the system is constructed. In a simple functional architecture, each component at a given level of partitioning performs roughly the same proportion of the total functionality as the other components at that level, and on each level the functionality is divided up between the components to minimize the amount of information which must be exchanged between components as they perform their functionality. With a simple functional architecture it is possible to relate functionality described at a high level to the device functionality with which it corresponds. Such relationships make it possible, for example, to identify the device problems associated with a failure experienced at a system level and take appropriate repair action.

Consider as an example a conceptual design process for a telecommunications switch. The highest level functional definition could be "handle the telephone service for a group of users which could number between 10 thousand and 100 thousand". Note that the functionality defined is implicitly an instruction. The next step is to imagine a partitioning of the system functionality into several subsystems and to define the functionality within each subsystem, which is then further partitioned. A possible initial partitioning could be that shown in figure 1. Each separation can be expressed as an instruction. The initial partitioning is tested by performing a series of scenarios such as the example in table 1 to determine the volume of information which needs to be passed between the subsystems as they perform their separate functions. Modifications to the partitioning are made to minimize this information volume while retaining roughly equal numbers of operations in each separation.

In a commercial electronic system, each functional component receives data which is meaningful and unambiguous to the component, and generates outputs which are instructions, or commands of specific system actions. Such a component can be regarded as detecting preprogrammed combinations of information, or patterns, which are unambiguous from a system point of view. One component can of course be programmed with many patterns. The "unambiguous" requirement is of critical importance, and can be better understood by an example. Imagine a simple calculator with two functional components, one of which receives two numbers from system input and multiplies them together, and the other displays an image based on the value of the product. The input to the first component is a pair of numbers, for example

33 and 26, which carry an unambiguous system meaning. Such information which is unambiguous at the system level is referred to as global information. Suppose that the first component multiplies the numbers by first calculating $33 \times 20 = 660$ and $33 \times 6 = 198$, then adding $660 + 198 = 858$. The number 858 could be communicated to the second component and would be unambiguous, the second component could generate an instruction to display the appropriate image. However, consider the numbers 660 and 198 generated inside the first functional component. Such numbers are unambiguous within that component, and could be shared between subcomponents of that component, but would be ambiguous if received by the second component. Such information is known as local information. If the local information 198 were communicated to the second component, it is not meaningless, such a value could only occur for a limited range of answers, it could not occur if the answer were 825 for example. The presence of 198 is therefore an indication of a range of possible images to be displayed, and a component receiving such information could in principle generate a recommendation to display a particular image, but would need correlation with other recommendations to create a high integrity system action.

In the partitioning in figure 1, the diagnostic subsystem cannot access information within call processing unless that information is expressed in a global system form. If the diagnostics subsystem as defined turns out to require high volumes of local call processing information, the effort to convert such information to globally meaningful form may be so large that is more effective to repartition the functionality between diagnostics and call processing. In the table 1 scenario, all the steps except the last would in the ideal case occur within the call processing function. An example of a global information requirement is for a database which relates all possible called numbers to a physical location. This database must be global because although it is used by call processing it is created and kept up to date by maintenance. Now consider the billing information in step 9. This billing information includes calling and called number, time of day and duration of call. If most calls are within a free dialing area and not billed to the customer, collecting this information and recording it in global form for every call would be an unnecessary drain on system resources. In such a case, it may be more effective to repartition functionality so that call processing makes the decision on whether to record the information, using a global database indicating destinations which are chargeable which would be kept up to date by billing. The process of design is thus one of rearranging the partitioning of functionality to achieve a tradeoff between even functionality and minimum information exchange. This design process is carried out at each more detailed level in turn until the device level is

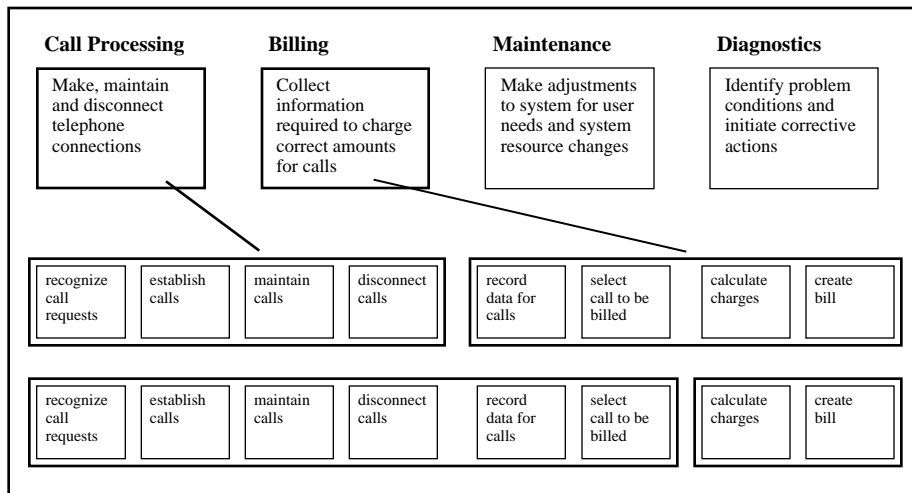


Figure 1 Initial high level partitioning of functionality for a telecommunications switch. The initial partitioning is tested using scenarios of the type illustrated in table 1 and modified to retain roughly equal functional components but minimizing the need for information exchange between components

reached. At any level, a necessary repartitioning may affect much higher levels, in the worst case a separation at the highest level may prove to be impractical at the device level, forcing repartitioning at every level. Design is thus a process of reaching a compromise between equal sized functional components

Scenario for a telephone call			
i.	User indicates to system the desire to make a call	vi.	System makes connection to called party
ii.	System indicates to user it is ready	vii.	System activated ringing of called party
iii.	User sends information on called party	viii.	If called party answers, system establishes link
iv.	System records information on called party	ix.	System checks periodically, and if one party disconnects the resources are reassigned
v.	System finds physical location of called party	x.	System records billing information

Table 1 A sequence of operations which could occur in the course of a system performing its functionality, used to test the information exchange required by a proposed functional partitioning

at each level and minimized exchange of information. At the most detailed level of a typical functional architecture, devices execute very simple instructions (open or close the device gate). A number of these device level instructions are combined to form assembly code instructions (jump, branch). Assembly code instructions are combined to form software instructions (while x is true, do y). Software instructions are combined into procedure calls, procedure calls into features, features into major system functions. This hierarchy is illustrated in figure 2.

There are several system consequences deriving from the partitioning of functionality into instruction components exchanging unambiguous information. Such a system requires an well defined source of such information: a memory function where a reference copy of unambiguous information is stored. While one component is active, it may be acting on some elements of unambiguous information, and the reference copy is therefore ambiguous for other components. Only one component can therefore be acting on global information at a time, and in general such a system has a processor function where the active components are executed sequentially. The only way to avoid sequential execution is if global information can be partitioned into completely separate and independent segments, and functional components can be defined which operate only with information in separate segments. Any system with the complexity requiring a functional architecture and in which functional components exchange unambiguous information is forced into an instruction architecture with its characteristic memory, processing separation and sequential execution. In most commercial electronic systems, much of the instruction functionality is coded as global information recorded in memory. This important instance of the instruction architecture is known as the von Neumann architecture. Because the functionality is treated as global information, heuristic modification of functionality, such as self modifying code, is difficult to implement.

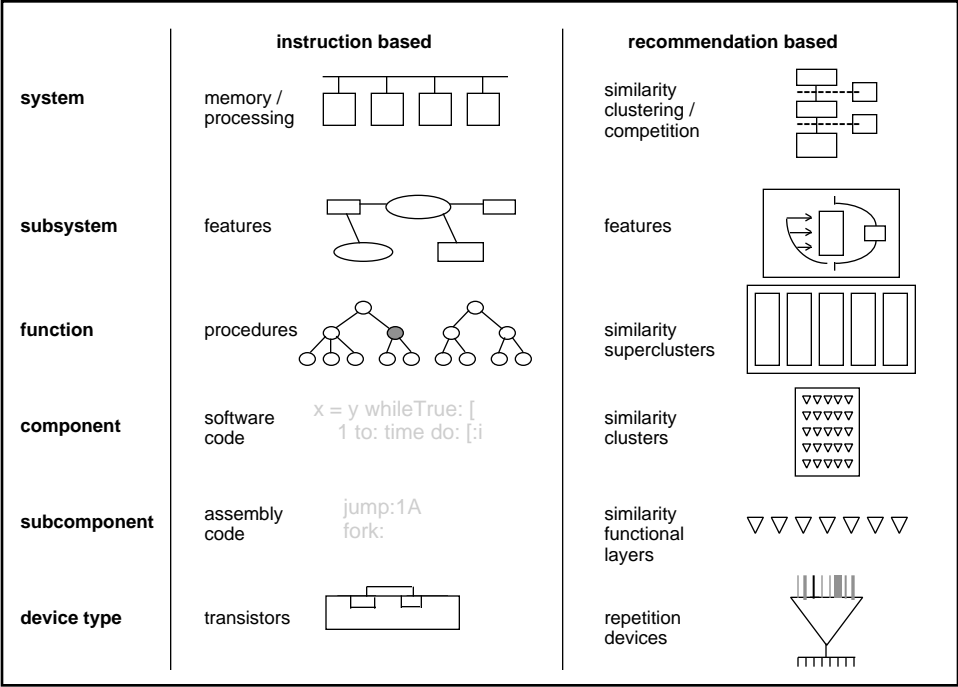


Figure 2 Functional hierarchies for alternative functional element paradigms. In the instruction architecture, instructions are combined into higher level instructions. In the recommendation architecture, repetitions are combined into similarity clusters. At each level a part of the higher level is shown, in the greater detail of the lower level.

If ambiguous information is exchanged between functional components, it is still possible to establish a simple functional architecture, but the form of that architecture is qualitatively different from that of the instruction architecture. The architecture allowing exchange of ambiguous information is the recommendation architecture (Coward 1990, 1997) and as discussed in Coward 1998 represents the only alternative to the instruction architecture for systems performing a complex functionality. When ambiguous information is the input to a functional component, the result generated by the component cannot be a system command, it can only be a recommendation which must compete with other recommendations for control of system action. The programmed combinations of data are not unambiguous system patterns but repetitions which can be extracted from the available ambiguous information and associated with action recommendations. In the functional hierarchy, repetitions detected at the device level are combined into clusters, which repeat if significant subsets of their component repetitions repeat. Superclusters repeat if significant subsets of their constituent clusters repeat. At the device level, the recommendations will be of the type "pay attention to this detected repetition" , at higher levels the recommendations will be increasingly general behavior types. The functional clustering hierarchy is illustrated in figure 2, compared with the instruction hierarchy. Note that if the information exchanged were unambiguous, the repetitions and clusters would be cognitive patterns and categories.

It is important to emphasize that because functional components exchange ambiguous information, the repetitions programmed at every level are not well defined cognitive patterns. The programming identifies the repetition of conditions which in many cases but not in every case indicate that an associated behavior is appropriate. An individual component is programmed to extract a wide range of repetitions with similar but ambiguous functional significance. The algorithmic complexity can therefore become considerable, even at the device level. For example, a functionally useful repetition could include the particular time sequence in which a combination of information appeared. The algorithms to detect such repetitions would be very complex.

Inputs to a component could in principle be information derived from any other component at any level. However, to maintain a simple functional architecture the distribution of information must be limited, and only made available where it has the highest probability of providing functional value. Optimization of information distribution is a critical system function, and has to be maintained by special purpose functionality in the heuristically defined functional architectures discussed in the next section.

The output of a cluster at any level indicates the presence of a condition under which the action recommended by the cluster may be appropriate. Different functional components at the same level acting on the same information may produce different recommendations. There must therefore be competition between the alternatives to determine system action. This competition requires some criteria to determine the recommendation which will be selected. The available criteria are derived from the results of accepting similar recommendations under similar conditions in the past. This knowledge could be utilized by design (i.e. for a biological system, by use of the results of past experience implicitly coded genetically) or heuristically from individual system experience.

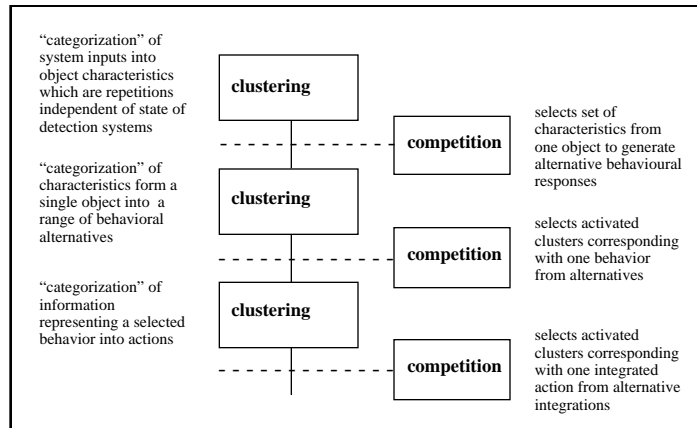


Figure 3 The major functional separations of the recommendation architecture are repetition similarity clustering of information and competition between the outputs from activated clusters for control of action.

Rather than the separation into memory and processing which occurs in the instruction architecture, in the recommendation architecture the major separation is therefore into repetition similarity clustering and competition. A biological example of these separations in such a system is illustrated in figure 3, following Coward 1990. Coward (1990) argued that, as shown in figure 3, in mammal brains there was a first stage of clustering which extracted sensory independent repetitions from raw sensory inputs. Cognitive examples of such repetitions might correlate with object color independent of illumination, or object size independent of distance. However, useful repetitions may or may not correlate with simple cognitive interpretations, they are only repetitions which can be discovered within the mass of ambiguous input information, and because the information is ambiguous would only correlate partially with such a cognitive pattern. The first stage of competition is between sets of repetitions extracted from different objects. This competition corresponds with the attention function, and the repetitions extracted from the object selected as the focus of attention gain access to the second clustering stage. This second clustering stage contains a number of superclusters corresponding with recommendations of different types of behavior (aggressive, friendly, food seeking, sexual etc.). The information derived from a single object may generate outputs from any or all of the superclusters, corresponding with different types of action recommendation with respect to the object. The specific combination of outputs from a supercluster indicate the specific action recommended within the type. The alternatives compete in a second competition stage and the information representing the successful recommendation is allowed access to the third clustering stage in which the output clusters correspond with portfolios of muscle movements. Competition between the outputs of this third stage results in an integrated physical movement.

Although biological systems have not gone through a design process, they are also subject to pressures towards a simple functional architecture. Copies of biological brains must be constructed from DNA "blueprints"; there must be some ability to recover from construction errors and damage; and random mutations must sometimes result in significant functional changes. There is therefore strong selection pressure in favor of simple neural functional architectures. Heuristically associating information repetitions with behaviors is fundamental to the existence of life. At the simplest level the detection of the repetition of a genetically programmed chemical gradient generates behavior in bacteria. At higher levels the detection of complex combinations of repetitions indicates the presence of a familiar person and generates a wide range of alternative behaviors. In a universe in which nothing ever repeated it is hard to imagine life dependent on learning from experience ever developing. Unambiguous information is typically not available to a biological brain, and Coward (1990) therefore argued that the recommendation architecture is ubiquitous in biological systems, and that the hierarchy of functional separations which therefore exists can be the basis for understanding psychological phenomena in terms of physiology. Evidence from biology is more fully discussed in Coward 1990 and in section 8.

3. Changes to Component Functionality in a System with the Recommendation Architecture

A component in a functional architecture receives input information and generates output results. There are two types of result (Mira et alii 1997). One type is any output information generated, and the other is any changes to the combinations of input information which will generate output information in the future. Changes to future functionality require input information recommending that changes occur, and such information will generally derive from a higher functional level than the component being changed.

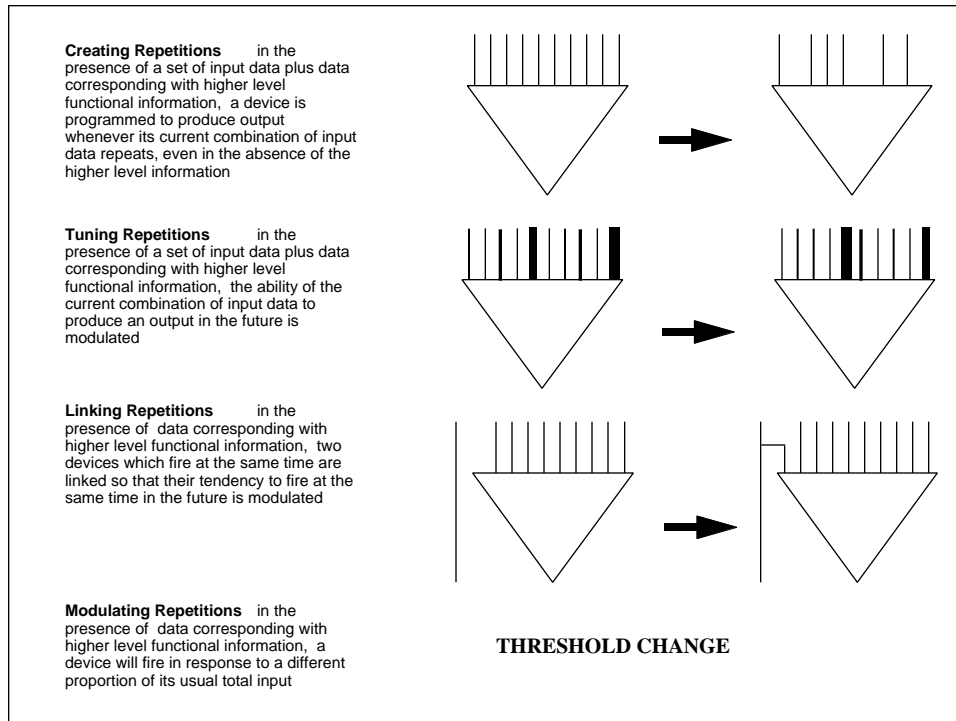


Figure 4 Alternative ways in which the combinations of local data to which a device is programmed to respond can change

At the device level, there are four different ways in which programming can change. These four ways are illustrated in figure 4. When a combination of input information is imprinted, the presence of information derived from a higher level component plus a set of input information causes a device to be programmed to produce an output at the time and at any point in the future if the same set of input information repeats even in the absence of the higher level information. In the simple example illustrated in figure 4, a device with a large number of inputs is imprinted by disconnection of all inactive inputs. The level at which the device will fire in the future is set at or below its current total input level. The nature and source of the higher level information is discussed below. This imprinting concept was introduced in Coward 1990. In its simplest form a device is programmed at a single point in time with a single combination, and produces output if that combination repeats in the future. The programmed combination will be referred to as a repetition if the combination is made up of ambiguous information, and pattern if the combination is made up of unambiguous information. Models with multiple repetitions per device are discussed below. Imprinting can support instantaneously created, permanent declarative memory traces (Coward 1990 and below).

An example of tuning of combinations is the standard perceptron model for the neuron (Rosenblatt 1961) in which the relative weights of different inputs are adjusted. An issue with many perceptron based neural network models, such as those based on backpropagation (Rumelhart et alii 1986), is the use of global information in the form of quantitative target answers. Use of such information will result either in an architecture which reduces to the sequential instruction architecture or a system which cannot scale up to handle very complex functionality. Tuning of combinations is important in the functional separations which competitively reduce the volume of data distributed to later clustering functions (Coward 1990). The competitive algorithms discussed by Taylor and Alavi (1993) are applicable within such separations.

Connecting combinations is a variation of the standard Hebbian mechanism (Hebb 1949) in which the connection between two devices which frequently fire at the same time is strengthened. The variation is that a connection may be established if two devices frequently fire at the same time. As discussed below, this mechanism is important in maintaining the orthogonality of multiple repetitions imprinted on the same device, and Coward (1990) proposed that the mechanism is also important in managing the distribution of information and in generating the self sustaining sequences of mental images which are one characteristic of human consciousness. All these types of combination changes depend at a more detailed level on mechanisms such as assignment or removal of devices; addition or deletion of single inputs; correlated addition or deletion of sets of inputs; changes in the relative strength of inputs; correlated changes in the strengths of sets of inputs; general changes in effective input strengths (i.e. threshold changes); and changes in the sensitivity to other parameters such as those which produce threshold changes.

Another qualitative aspect to changes in future functionality is the permanence of the changes, which can vary from change only while the controlling higher level functional information is present through limited time to long term after the controlling information is no longer present.

The particular combination of functionality changes appropriate to a device is determined by the higher level functional component within which the device is located.

4. Heuristically Generated Functionality in the Recommendation Architecture

In a system with the recommendation architecture, a cluster hierarchy can be established heuristically as a simple hierarchy of repetition. At all levels there is a search for repetition of sets or large subsets of information combinations which have occurred before. Frequent repetition results in the repetition being established as a component of similarity at its appropriate hierarchical level. Raw input information is searched for combinations which repeat, and any potentially useful repetitions are programmed into devices. These combinations are searched for combinations of combinations which repeat, and such combinations of combinations are programmed into clusters. Frequently repeating combinations of clusters are combined into superclusters and so on. The resulting similarity hierarchy can then be heuristically associated with behaviors. A similarity cluster can be provisionally associated with a particular type of behavioral recommendation, and if the results of accepting the recommendation are favorable, the association can be made permanent, otherwise it can be discontinued.

The repetition imprinting mechanism thus makes it possible to define functionality heuristically on the basis of experience. The process depends upon maintaining a clear separation between heuristic definition of clusters and heuristic association between clusters and appropriate behaviors, plus a clear functional hierarchy and management of information flow across that hierarchy.

In figure 5, a cluster module composed of devices which can be imprinted with repetitions is illustrated conceptually. The cluster module is made up of a succession of layers of devices. As discussed later these layers are functional separations and in general the number of layers depends on the number of functional separations needed. A device has inputs corresponding with elements of information which are outputs from other devices, and is activated if a large enough proportion of its inputs are activated. Most inputs to devices within the module are from other devices in the module, and a high proportion of the inputs are from devices in the preceding layer. Lateral inputs and back projection inputs are also possible, but these inputs represent information from higher functional components, in this case layers. The first layer receives inputs from outside the module, and such inputs may be received by other layers for functional reasons such as maintaining information consistency as discussed later. Regular devices have already been imprinted with their repetition, virgin devices have not yet been imprinted, and have more inputs than regular devices, but drawn from a similar population to the inputs to regular devices in the same layer.

Higher functional connectivity in the illustrated module is of two types. The first type is feedback from regular and virgin devices in the output layer to virgin devices in the output and all earlier layers of the cluster. This feedback is from a high proportion of output devices to all virgin devices, and inhibits imprinting. The second type is from regular devices in the middle layer to virgin devices in the same and all other layers, there are again a large number of such connections and the connections stimulate imprinting. To avoid confusion, examples of feedback connectivity are portrayed rather than a realistic set as described in the text.

Suppose that the module in figure 5 has already been programmed to generate output in the presence of one (ambiguous) condition or object. Such programming would mean that input level regular devices are programmed with combinations of information extracted from the condition, regular devices in the second layer with combinations of input layer devices and so on, with the result that if information from the same

condition were presented again, an output would result. Because two perceptions of the same object will not necessarily be identical, actual device activation would be a large subset of the total, but an output from the module would be generated. The inhibitive inputs to virgin neurons from those outputs dominates, and no imprinting would occur.

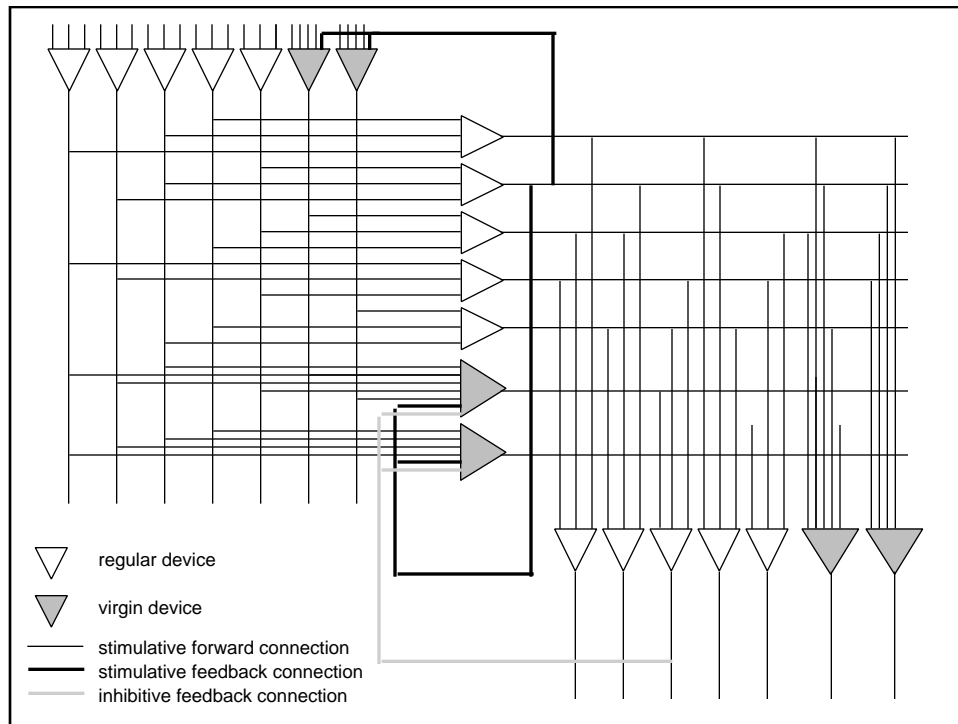


Figure 5. Connectivity of a simple repetition similarity cluster module using the repetition imprinting mechanism at the device level. The layers perform similarity subfunctions as discussed in the text. Single examples of connectivity which performs required layer to device functionality are given, realistic functionality requires many more connections as discussed in the text

Now consider what would happen if another condition very similar to the first were perceived. Many of the simple combinations of information in the early layers would repeat, but the very complex combinations close to output would not repeat. The stimulatory inputs to virgin devices would therefore be present, but not the inhibitive inputs from the output layer. Under these conditions, the number of active inputs from the preceding layer required to fire virgin devices is gradually lowered (for example, in proportion to the length of time the stimulate feedback has been present), and these devices begin to fire. The thresholds drop until enough virgin devices fire in the different layers to generate an output, at which point further imprinting is cut off by the inhibitive feedback. At first firing, a virgin device is imprinted so that any future repetition of a large proportion of the information combination will cause the device to fire, independent of the various feedback mechanisms. The new condition will therefore generate a module output whenever it is present in the future. Note also that the new repetitions imprinted by the new condition will contribute to triggering imprinting in future objects, and the definition of the cluster is thus generated heuristically. Because the input information is ambiguous, the outputs are ambiguous, in other words, because the recurrence of the same condition may not result in exactly the same inputs, and because other conditions may result in inputs similar to the original inputs, the output ambiguously correlates with the cognitive condition.

A condition which was not similar to the cluster definition would generate limited firing in the early layers, but insufficient to generate imprinting, and such a condition would have no effect on the cluster. A new cluster would be initiated by the experience of conditions which produced no response from any existing cluster. Devices in such a new cluster would be configured with large numbers of randomly selected inputs of appropriate types. The first condition which subsequently produced no response from

existing clusters would force imprinting in the new cluster. The new cluster would then heuristically develop an internal similarity definition by experience of similar conditions.

One module being presented with information from various conditions is shown in figure 6. The connectivity is omitted, and five layers which are assigned different functional roles in later discussion are illustrated. For ease of explanation, initially a cognitive example implying the use of unambiguous information will be described and then the differences with ambiguous information explained. Suppose that the functional cluster component is a category (i.e. unambiguous) associated with actions of the type “say that the painting is by Magritte”. At the upper levels of the module, the devices detect patterns of relatively simple combinations of sensory data. The middle layers detect patterns of complex combinations of information which correspond with the 'style' of the painter. Close to output, very complex combinations are extracted which in general are unique to individual paintings. The exact combination of outputs can therefore be used to control the painting specific content of the action. If a painting by another artist is presented, there may be a limited extraction of patterns at the input levels and such weak extraction is ignored. If a new painting by Magritte is presented, there will be strong detection of patterns reaching deep into the module, but no output will be generated. Under these conditions, additional patterns are permanently imprinted until an output results. The new painting will be recognized in any future presentation because no significant imprinting will be required to produce an output. The additional patterns are additions to the category definition because they can contribute to triggering imprinting in response to future paintings.

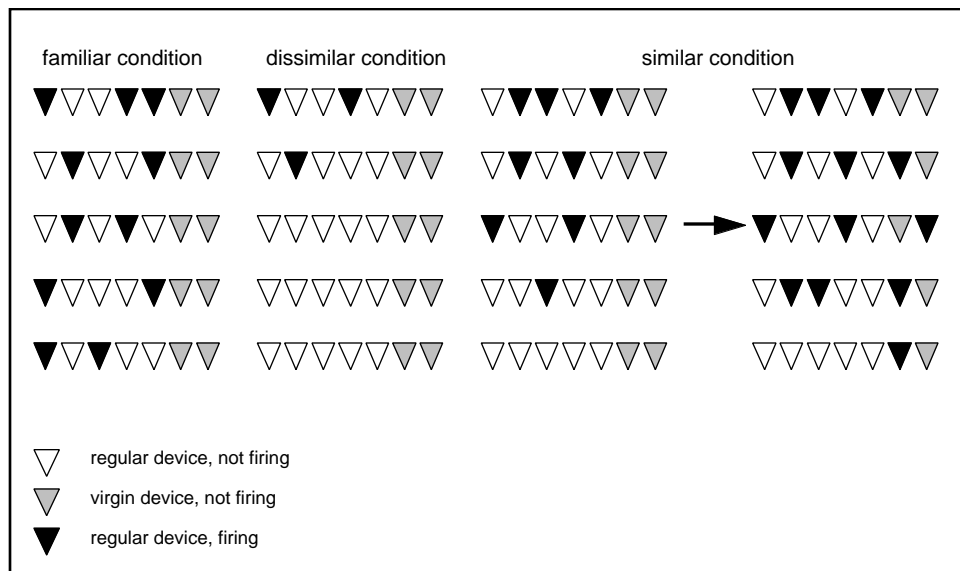


Figure 6 Activation of one cluster module in response to information extracted from a condition which has been experienced before, a condition which is dissimilar from the objects which have defined the cluster in the past, a condition similar to past conditions at the moment the information is initially presented, and after imprinting. Devices in one layer are activated by programmed combinations of information, largely derived from activation of devices in the previous layer

Now consider the process with limitation to ambiguous information. Suppose the system is presented with a series of sensory experiences include experiences of a number of different painters. These experiences trigger the creation of a number of clusters corresponding with heuristically defined repetition conditions. These repetition conditions do not correspond with individual painters, but indicate the presence of similarity conditions in input information. Several clusters might produce output in response to a particular painter, not always all of the same clusters for a particular painting by that painter. Paintings by a different painter might produce outputs in different clusters, but sometimes the same cluster might produce outputs for different painters. The output of these clusters is therefore ambiguous, but can be used by a competitive function to generate high integrity behavior. For example, suppose there were five painters A through E, and the heuristic process generated seven similarity clusters one through seven. A recommendation to *say the name of painter A* could become associated by trial and error with strong

outputs from cluster two, weak outputs from cluster four, and moderate outputs from cluster seven, a recommendation to *say the name of painter B* could become associated by trial and error with strong outputs from cluster four, weak outputs from cluster five, and moderate outputs from cluster six etc. A simple competitive algorithm could be designed which after an initial trial period would result in any novel painting by one of the five generating the appropriate response.

There are a number of interesting properties of this type of clustering process (see Coward 1990). One property is that all the repetitions active at the time a condition is perceived are permanently programmed. If these combinations could somehow be activated again, the experience would be indistinguishable from the original exposure. Secondly, a permanent trace is instantaneously created which permits a clear distinction between conditions which have been experienced before and other conditions. Conditions not experienced before will require significant imprinting to generate an output, while conditions seen before will require no or much less imprinting. Thirdly, because the combinations are simply repetitions of information which was present at the same time, the repetitions associated with one condition may include information extracted from another condition which happened to be present. The infrastructure to support associative memory is thus established as a side effect of the clustering process.

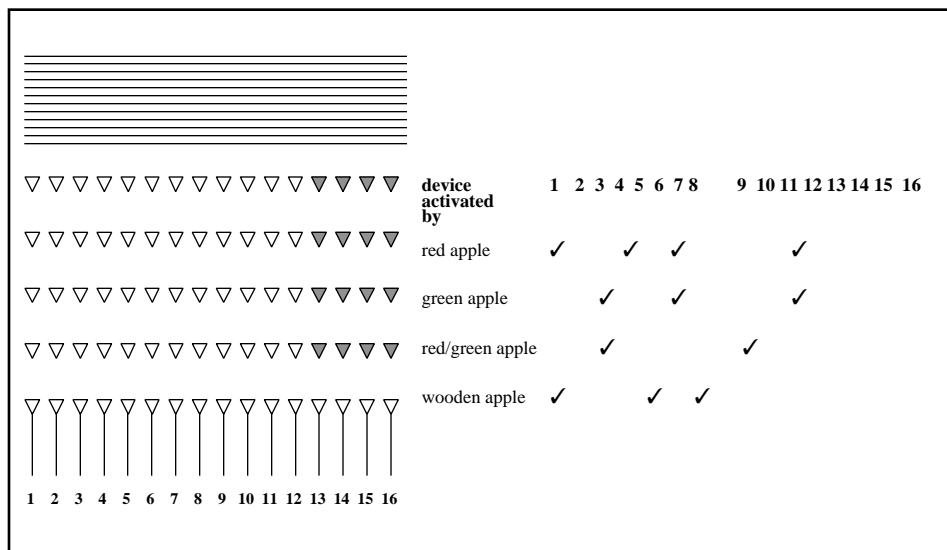


Figure 7 Ambiguous information output by a cluster. Any output indicates the presence of a condition similar to cluster definition, and can be used as an action recommendation. The illustrated cluster could recommend eating behavior. The particular combination of output information is generally unique to a particular condition. The uniqueness can be used by the subsequent competition to manage the probability of future success of accepted recommendations.

The process for establishing an heuristic association between clusters and behavior can be understood in more detail by consideration of figure 7. Suppose that the illustrated cluster module is one of several which have heuristically acquired the ability to generate an output in response to apple like objects. In principle it can be imagined that the outputs from such clusters could be distributed to functions driving all types of behavior. Suppose that one type of behavior were eating. Output from the module will then be interpreted as an action recommendation to eat. Now suppose that other clusters exist which detect biologically valuable and biologically unfavorable conditions. The output of a biologically valuable condition cluster is a recommendation to strengthen the connections from the clusters which recently gained control of behavior on to the recently active behavior type, the output of biologically unfavorable condition clusters is a recommendation to weaken such connections or in extreme cases to disconnect any outputs from the source clusters.

An 'apple' cluster will generate outputs in response to many apple like objects, including apples of different colors, and even wooden apples and other objects with the right type of similarity to apples. However, the particular combination of outputs are unique to individual objects. Hence the output of the clusters indicating the presence of biologically favorable or unfavorable conditions will result in disconnection of recommendations in response to wooden apples while enhancing the strength of

recommendations in response to real apples. A module which has heuristically acquired the ability to generate outputs in response to generally inedible objects would be disconnected from any access to eating behavior after experimentation. In system terms the repetition similarity condition detected by the cluster is not functionally valuable for recommending such a behavior. More complex competitive algorithms which can use outputs to modulate behavior within the general type recommended by the cluster are discussed below. Once the output from similarity clusters has been associated through a competitive function to a behavior, the behavior will be associated with a cognitive category, although the operational definition of the category is the weighted average of the presence or absence of a range of ambiguous repetition similarity clusters.

Electronic simulation has demonstrated that using the algorithms described in this section a system with the recommendation architecture can organize its experience into clusters without guidance, and the clusters can generate high integrity behavior through a competitive function after a small number of trials using only correct/incorrect feedback (Coward 1996, 1998).

5. Similarity Hierarchy in a System with the Recommendation Architecture

The next higher level of the similarity clustering functionality can be understood by consideration of figure 8. All perceived conditions (or objects) are sorted into clusters, with individual conditions being assigned to the clusters which generate the strongest familiarity indication as illustrated in figure 8. Initiation of a cluster is triggered by a single condition which does not generate output from any existing cluster. A new cluster module with large, random connectivity between device layers is forced to produce outputs in response to such conditions. Any subsequent condition which is sufficiently similar to the initial condition will generate imprinting in the same cluster, and the additional repetitions will refine the cluster definition. Correct/incorrect feedback can be used to guide the association between clusters and recommended behaviors. This functional separation is a key aspect of the architecture.

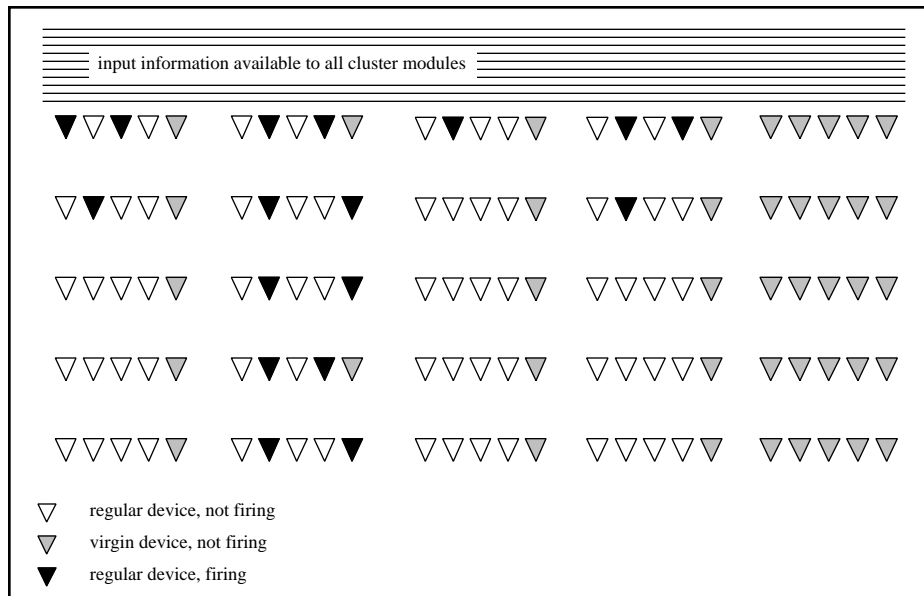


Figure 8. Information repetitions extracted from a condition or object are presented to a range of clusters, and the clusters with the strongest activation imprint additional repetitions to produce an output. Conditions are thus heuristically sorted into clusters.

Note once again that if the information were unambiguous, the clusters would be categories. Clusters are ambiguous, and will therefore correlate partially but not exactly with cognitive categories. Neural network algorithms can achieve separation of input into categories. For example, Kohonen maps have demonstrated the ability to self organize to extract a reasonably small set of important features from input information (Ritter 1995). However, Kohonen maps target unambiguous information as output, and even Kohonen's topological map scheme for unsupervised learning has been criticized on the basis that the patterns which are presented to the model must be preselected (Pfeiffer 1996). Furthermore, gradual

adjustment of weights has difficulty in accounting for the instantaneously created ability to recognize that an object has been seen before as discussed earlier.

The primary issue with any similarity process for functional purposes is whether it can converge on a useful set of similarity clusters or categories. There are two problems to avoid. One is a proliferation of clusters or categories which at the extreme could reach one per object or even more. The second is the combining of many objects into a single category when the appropriate behaviors are not the same. Because there is scope for correction of the behavioral consequences of errors through the separate competition process as discussed, it is not necessary that the clustering process even in principle be capable of perfect separation into cognitive categories, the requirement is that the resulting repetition similarity separation be practically useful. To put the same statement another way, a functional architecture made up of functional modules exchanging ambiguous information is able to generate appropriate actions. A further point to emphasize is that clean separation between clustering and competition is essential to achieve a simple functional architecture. The lack of this clean separation in neural networks is one reason that they cannot utilize ambiguous information in complex functional combinations.

Although the clustering can proceed without guidance, there are a number of parameters which must be defined, and feedback of the effectiveness of the clustering process can be used to tune the value of these parameters. These parameters include the number and configuration of virgin devices available for imprinting of useful repetitions; the balance between inhibition of imprinting by output and lateral stimulation; defining the level of lateral stimulation which will trigger imprinting; and duplication and proliferation of clusters. The critical issue is to ensure that the process of heuristic similarity clustering achieves a cluster separation which is useful in generating behavioral alternatives.

Feedback can be used to tune the parameters in a number of ways. For example, if a cluster produces recommendations which if accepted are sometimes followed by favorable conditions, sometimes unfavorable, a system management cluster detecting the conflicting results could generate recommendations to tighten the similarity definition, leading to more clusters. If multiple clusters frequently generated recommendations of the same type in response to the same object, other system management clusters could generate recommendations to merge the duplicate clusters. If within a cluster an object triggered imprinting but there were not enough combinations of the right type present to generate an output, this condition could trigger configuration of more or larger combinations. A tendency to proliferate clusters could be reduced by adjustment to the criterion for imprinting within clusters.

Within one heuristically defined cluster, the new imprinted repetitions required for adding a new condition to the set which produce cluster output are combinations of active inputs to individual devices prepared to accept new repetitions. The effectiveness of learning will be influenced by the available combinations of physical inputs on such devices. These physical combinations must in general be configured in advance in a system dependent on physically wired connections. The configuration process could be random, but because useful new repetitions will frequently include new combinations of information elements which frequently occur in other repetitions in the same cluster, a bias in favor of inputs corresponding with such information will improve the effectiveness of learning. Such a bias could be achieved by activating all devices in proportion to the frequency of past activations and causing devices which will be programmed with additional repetitions to accept inputs which frequently contribute to firing other devices in their functional neighborhood. Such an activation would resemble a fast rerun of a mixture of past experience, and Coward (1990) proposed that providing this type of bias on connectivity was a primary functional role of dream sleep in biological brains.

This management of the information available at the device level is a specific example of an important architectural point. Although in principle a functional component at any level could accept input from any other component, the functional architecture is more simple if the information exchange routes are minimized. A compromise will be reached between the need for a simple functional architecture and the need for the most appropriate information to be available to each functional component. There is evidence that in biological brains the compromise is that each functional module at the cluster level receives input from about ten other modules and delivers output to about ten modules (Bressler 1994). Selection of the information which will be used as input to a cluster is the function of the cluster input layer. Provisional inputs from other clusters would be accepted if those cluster outputs were frequently active at the same time as the receiving cluster was also generating outputs, provided the coincident outputs were more frequent overall than with other clusters. Dream sleep manages this information distribution function in a manner analogous with the device level information management described above.

A number of separate cluster subfunctions have been identified in the above discussion: managing inputs from other clusters; detecting the presence of enough activation to inhibit the creation of a new cluster; detecting the presence of enough activation to stimulate addition of the condition to the set which produce cluster output; detection of enough feedback from higher clusters to maintain activation; detection of enough output to inhibit additional imprinting; and generation of unique condition outputs. These separate functions can be separated in different module layers. The advantage of such a separation is that the functionality of one layer can be adjusted without introducing uncontrollable changes in the functionality of other layers.

Note that at every functional level, the clusters and repetitions which are programmed are simply combinations of information which have occurred once, and become potentially useful if a large subset of the combination occurs again.

Any domain of experience will be similarity clustered in order to associate different behavioral recommendations with different cluster combinations. Even continuously variable experience will be broken up into clusters in a system with the recommendation architecture. The experience of the forced separation of the continuous distribution of light wavelengths into the seven colors of the rainbow is an illustration of this effect.

At a higher level in the functional hierarchy illustrated in figure 9, superclusters generate alternative general types of action recommendations. The superclusters correspond with alternatives such as aggressive, fearful, etc. Each supercluster contains a set of cluster modules to generate recommendations of the supercluster type towards any possible condition. Each of the active combinations in figure 9 could be regarded as “dog recognizing”. A more accurate interpretation is that the different modules generate different action recommendations with respect to a similar but not necessarily cognitively identical "dog" cluster set. The supercluster separation also allows modulation of the different types of action recommendation depending

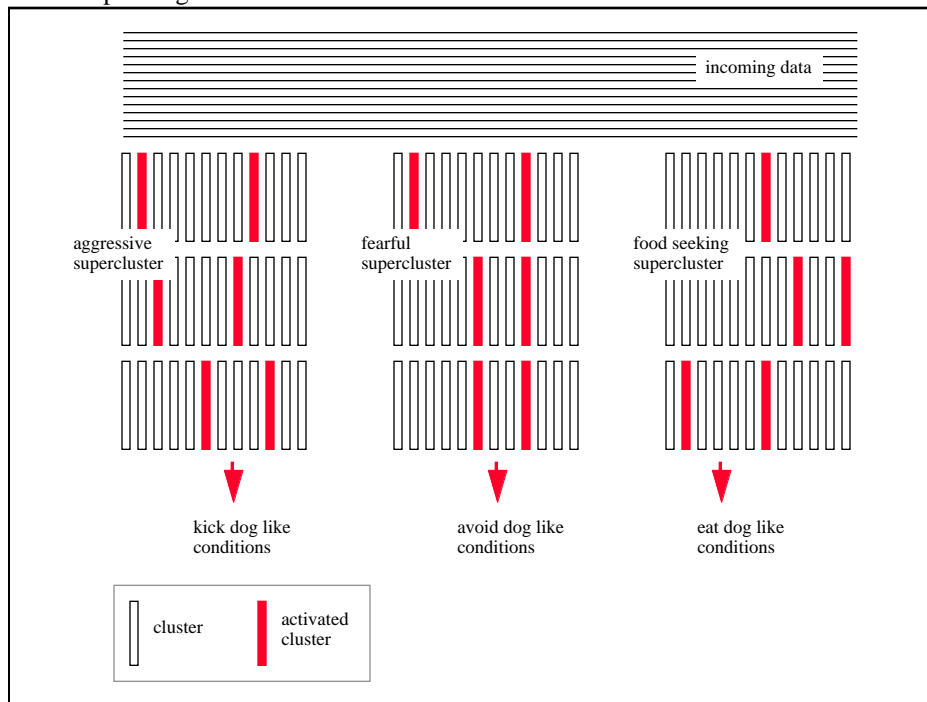


Figure 9 Parallel hierarchies of heuristically created clusters generate alternative behavioral recommendations towards the same perceived object or condition. The activation of a set of clusters corresponds with the recommendation.

on information about system needs. For example, extraction from internal sensory data of a cluster set which could be cognitively labeled “hunger condition” could correspond with an action recommendation to reduce thresholds of all devices in the food seeking supercluster. The duplication of clusters which could all be interpreted as “dog recognizing” means that local damage to the system cannot remove all the memory

trace of a particular dog experience. The effect of damage would rather be to change the probable response of the system towards dogs. Such duplication could in principle be regarded as wasteful, but represents the solution to the requirement for a clear functional separation under conditions of ambiguous information usage.

The simple imprinting algorithm used implicitly in the above discussion can be modified to allow devices to be programmed to extract multiple repetitions imprinted at different times. Such a modification must solve the orthogonality problem illustrated in figure 10. How can a target device C1 in layer C which has been imprinted with a set of repetitions including B1(p1) in layer B be prevented from responding to the activation of B1(p2). The simplest solution is to require that multiple repetitions on a device like B1 do not have any inputs in common, and that their targets like C1 also receive some inputs from layer A which is the primary source for inputs to B1. This structure is still feedforward, but net repetition complexity increases somewhat more slowly from layer to layer than in the simple case. In practice this approach would probably sustain a practical degree of orthogonality even if there were a small degree of input sharing between different repetitions on the same device.




Figure 10. Repetition orthogonality problem and solutions. If a device is programmed with multiple repetitions, its target devices must be able to distinguish between outputs due to relevant repetitions and irrelevant repetitions. One mechanism is to forbid sharing of inputs between different repetitions on the same source device and allow some of the inputs to the source to reach the target. A less restrictive mechanism is to use feedback to converge on a consistent set of device activations

If the sharing becomes significant, an additional mechanism is feedback connections from C1 back to earlier layers such as A. Such feedback connections could be established at the time of imprinting of a repetition on C, and have the effect that an output from the recipient of the feedback would only continue if significant feedback were received. Such a mechanism would drive convergence on a consistent set of repetitions through many layers, the consistency being that the set have tended to be extracted at the same time in the past. In practice such feedback would be more useful if it extended from the more complex clusters to the simpler clusters illustrated in figure 9. Cauller (1997) has proposed that the extensive feedback connectivity observed in biological brains has a predictive role and drives convergence of a dynamic system towards an attractor. The mechanism described here is functionally equivalent but expressed in recommendation architecture terms. Because a complex system with many components must employ either recommendation or instruction functional partitioning, understanding of a system activation will be most accessible in the applicable functional paradigm. The tendency of the feedback circuitry described by Cauller to target a particular cortex layer reflects the same requirement for functional separation within clusters discussed earlier.

The feedback mechanism allows simple separation of the functionality to maintain orthogonality between the multiple repetitions imprinted on a single device. Higher level components can therefore use either the simple or more complex device functionality to achieve the same higher level functionality. The simple mechanism can therefore be used to understand the higher level functionality in device terms, the separate orthogonality mechanism can be added if required for system operating efficiency reasons.

The repetition similarity clustering hierarchy which results from the discussion of this section is illustrated in figure 11, using a biological example for illustration purposes. At the highest functional level illustrated, the superclusters correspond with broad behavioral types, and an output from such a cluster is a recommendation of a behavior of that type. Such an output is composed of outputs from a set of clusters within the supercluster. Such a set can be interpreted as an indication that an object or condition similar to a past condition is present and a similar behavior to one tried in the past is therefore recommended. The specific outputs can be used to specify the particular type of behavior within the general type. Within a cluster, layer subcomponents perform similarity subfunctions which are in turn made up of device level functions. An important architectural point is that if a set of information can be defined which is separate from the primary information set and can be used to generate a separate set of action recommendations, then it will be functionally simpler to establish a separate clustering, competition subarchitecture. Two such subarchitectures are illustrated in figure 11. One clusters information indicating system needs (e.g. hunger) and generates action recommendations to modulate the relative probability of a type of action (e.g. food seeking) by lowering the thresholds of a set of devices within the appropriate supercluster. Such a result could be achieved by distribution of an arousal neurohormone targeted at the supercluster (Coward 1990). The other subarchitecture clusters information deriving from the effectiveness of the clustering process and individual cluster activity and generates resource assignment recommendations including assignment of additional resources to clusters demonstrating high imprinting rates. Such a function would contain implicit information about the time sequence of experience which could be used for behavior recommendation generation.

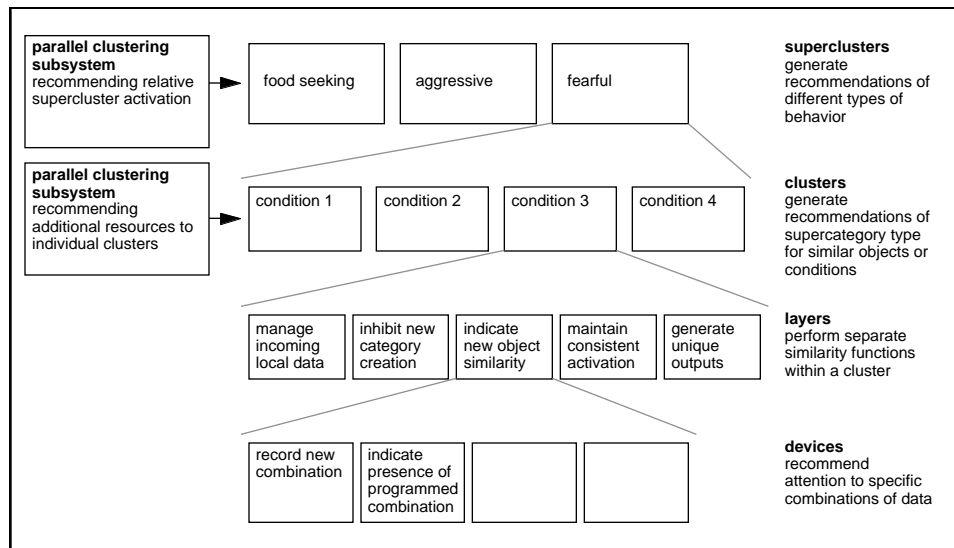


Figure 11 A repetition similarity clustering hierarchy. Separate subsystems can exist with their own clustering hierarchy generating recommendations for management actions on the primary hierarchy. Separate competition functions will manage the access of these recommendations to the primary hierarchy.

If a totally novel type of condition were perceived, then new clusters would be established in all behavioral superclusters. Devices in the new clusters would be made sensitive to the appropriate arousal neurohormone. The information used as input to the clusters would be biased towards information which has generally proved useful in the past in generating behavior of the supercluster type. Behavior towards such a totally novel condition would be experimental, with many different types being tried. In a mature brain, most conditions would have some similarity to past experience and could generally be accommodated with an additional cluster in an existing set. It is also possible for a new supercluster to be defined heuristically. Such a supercluster would have a different general arousal sensitivity, perhaps a sensitivity to a combination of primary neurohormones. Clusters would be heuristically established within the supercluster. An example could be game playing or sport, where neurohormone sensitivity could include combinations of cooperative and aggressive and other sensitivities. For a more complete discussion see Coward 1990.

Within the similarity clustering based functional architecture, memory as experienced by the system has a number of meanings. Firstly, if an object is perceived and little imprinting is required to generate an action recommendation, the object will be experienced as familiar. If significant imprinting is required to generate action recommendations, the object will be experienced as unfamiliar. Perception of the degree of imprinting required to generate behavioral recommendations is thus the device level basis for the instantaneous creation of a permanent trace giving the ability to recognize familiarity. Although there is some resemblance between the mechanism for similarity triggered learning described here and the adaptive resonance algorithms (Carpenter and Grossberg 1988) there are substantial differences. ART tunes patterns using correlated firing of devices in the neighborhood. The imprinting mechanism (Coward 1990) uses similar information to manage the creation of repetitions. A key difference between the imprinting algorithm and adaptive resonance is that a single exposure to an object results in a permanent record, and a clear difference is established between the experience of an object which has been seen before and an object which is novel. A further difference is that ART targets output of unambiguous category vectors and ART modules therefore cannot be combined in a recommendation architecture.

Secondly, the immediate experience of a perceived object is the activation configuration generated, made up of device activation across the alternative recommendation type superclusters. This activation can be imagined as a set of parallel cascades of device firing, corresponding with alternative behavioral recommendations. As long as a path from input to behavioral recommendations is activated, or is in an aroused state for which a small proportion of original input will reactivate, the recommendations can be regenerated and the object experience maintained. The aroused state is the basis for short term memory, during which the partial mental state induced by an object can be maintained.

Thirdly, the repetition recording mechanism at the device level imprints the currently active combination of inputs. Actual imprinted repetitions will therefore depend upon whatever physical inputs happen to be available. Even with the biasing provided by dream sleep some physical inputs to devices in clusters activated by the perception of a dog would be activated by the perception of, say, a cat or a tree. If such objects happened to be present when new information combinations are imprinted by a novel dog, then some information from such extraneous objects could be included. As an example, the perception of a dog chasing a cat up a tree could result in some cat and tree information being included in repetitions imprinted in dog related clusters. Some of the set of cat related clusters could even be shared with the set of dog related clusters. A later perception of a dog and a tree could result in a weak associative activation of a cat related behavioral recommendation. This activation is the basis for associative memory, which is the way in which images and behaviors appropriate for objects which are not present can be generated. Unless amplified in some way, such a weak associative activation would be negligible compared with the primary activations. The means by which such an amplification could occur is discussed later.

Fourthly, the use of the imprinting mechanism means that within the similarity clustering function all the information combinations activated at the time of an experience are permanently recorded, although only a small subset are actually recorded during the experience, the majority are repetitions of combinations recorded in earlier experiences. Hence if associative activation were able to activate the complete set, the experience would be indistinguishable from the original perception. In practice as discussed below, in reminiscence the activation is always of a subset which in particular excludes repetitions close to sensory input. One qualification on this argument is that if repetitions which never or extremely rarely repeated were eventually deleted to save resources and simplify the functional architecture as proposed in Coward 1990, then such repetitions would no longer be available.

Because similarity clustering is heuristic, and therefore strongly influenced by the type and order of experience, it will vary between different individuals, and individual behavior will be strongly influenced by individual experience. Furthermore, the repetitions programmed in an individual brain are determined partly by what connectivity options happened to be available. Hence the activations in two brains in response to the same experience will vary both because of differences in experience and because of accidental connectivity factors. Experience of a specific external input, say the color blue, is the activation of an assembly of repetitions created in the course of a wide range of individual experiences of the color blue in conjunction with other experiences. These repetitions are unique to individual experience, and the experience *blue* is therefore unique to the individual.

Different types of component output information can be the basis for an additional functional separation. For example, the presence or absence of an output can indicate the presence or absence of a recommendation, the modulation of that output (e.g. device firing rate) could indicate the strength of the recommendation, and the second order modulation (e.g. rate of change of device firing rate) could indicate

the input data population from which the recommendation derived. The second order modulation could thus be the functional basis for object binding as proposed by Llinas et alii (1994).

6. Competitive Reduction of Local Information

The above discussion of heuristic association between clusters and behavior included a simple mechanism by which information output by a cluster could be eliminated from consideration by a management mechanism using simple pleasure and pain type feedback. A practical architecture requires more sophisticated means to reduce the volume of information. Similarity clustering of input information can generate a large volume of ambiguous information corresponding with action recommendations of different types. In the early stages these recommendations can be interpreted as "pay attention", at later stages interpretations are more specific system actions. In the multistage reduction suggested as a model for the mammal brain in Coward 1990, the first stage is a competition between domain outlines which are ambiguous equivalents of the primal sketches described by Marr (1982). This mechanism is more completely discussed in Coward 1997a, and the result of the competition is access for the information originating within the domain defined by the primal sketch to more detailed clustering generating alternative behavioral recommendations. The second competitive stage as described in Coward 1990 is illustrated in figure 12. Recommendations from different superclusters enter different but parallel pipes, with cross inhibition between pipes which is modulated to achieve the exit of not more than one recommendation from the structure. The success of a recommendation allows the corresponding information access to a further clustering into body movements. Coward (1990) suggested that the first competition corresponds with the thalamus structures, the second with the basal ganglia, and offered supporting physiological evidence.

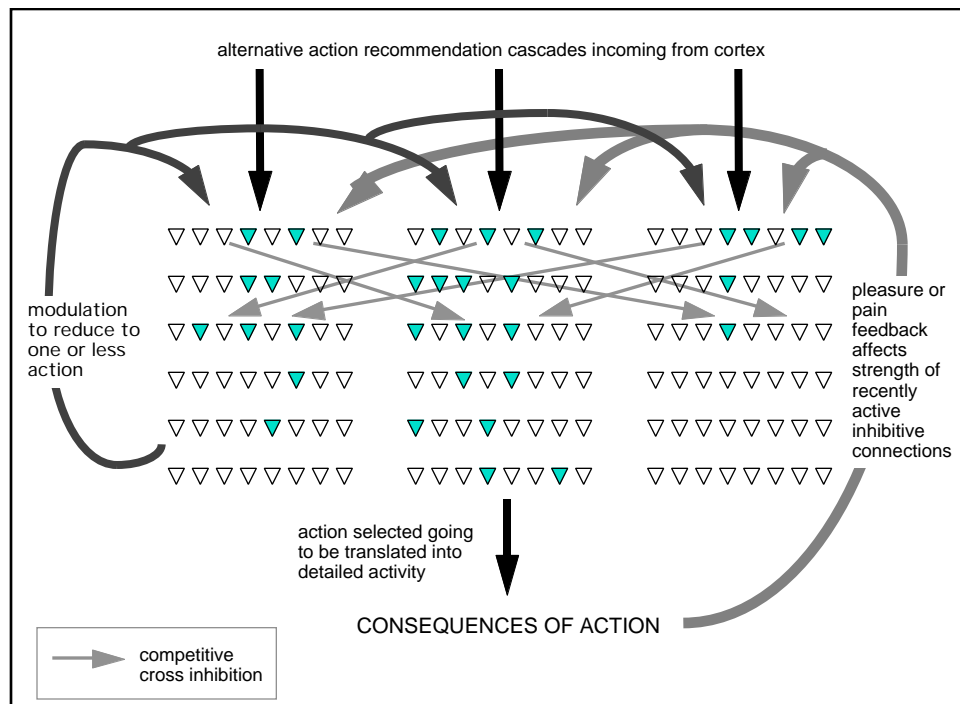


Figure 12 A possible competition structure. Active information representing alternative action recommendations of different types enters separate pipe like paths. The recommendation in one path inhibits the transmission in all other paths. If more than one recommendation approaches output, a feedback loop damps down all paths. Feedback on the consequences of an accepted recommendation changes the strength of recently active inhibitive connections in the successful path and thus modulates the probability of similar actions in the future

7. Cluster Management in a System with the Recommendation Architecture

The process of imprinting of repetitions at the device level can support a range of cluster management functions in addition to the basic similarity clustering of experience described earlier. For example, once the cluster structure to generate a set of behavioral recommendations has been heuristically defined, there may be advantages in accelerating the speed with which a response can be generated in the presence of its programmed information. Such an acceleration could be achieved by reducing the number of device layers through which information must pass in the cluster hierarchy, at a cost of reducing its ability to evolve heuristically in the future. Such a reduction could be achieved by setting up a parallel but simpler cluster structure, providing that structure with the same input and output populations as the primary structure, imprinting paths through the simpler structure only when the primary structure was producing outputs, and imprinting the outputs of the secondary structure to the targets of those primary outputs.

The process described is essentially one of cloning an existing cluster structure for the same function. It is also possible to clone a cluster structure for a different function. For example, as described in Coward 1990, consider a system which had already established the cluster structure to play table tennis, which is required to acquire the capability to play tennis. At some level of clustering there is some similarity between the two activities. If the clusters developed for table tennis are forced to generate behavioral recommendations for tennis, partially usable recommendations will result, with imprinting of additional repetitions. The new repetitions will incorporate information more appropriate for tennis, and the outputs will become heuristically associated with behaviors more appropriate for tennis. The newly imprinted repetitions will be the seeds of clusters fully appropriate to tennis. Creation of a cluster hierarchy from scratch would require a much longer and more resource intensive process.

Another application of the cloning process is suggested in figure 13. If clusters have developed which generate action recommendations appropriate to certain shapes being in particular relative position, their inputs will be information relevant to the presence of the shapes and their relative positions. Now suppose that information on the relative time sequence of perception of the objects is available, for example from the hippocampus as a result of the resource assignment process suggested in Coward 1990. If simultaneous

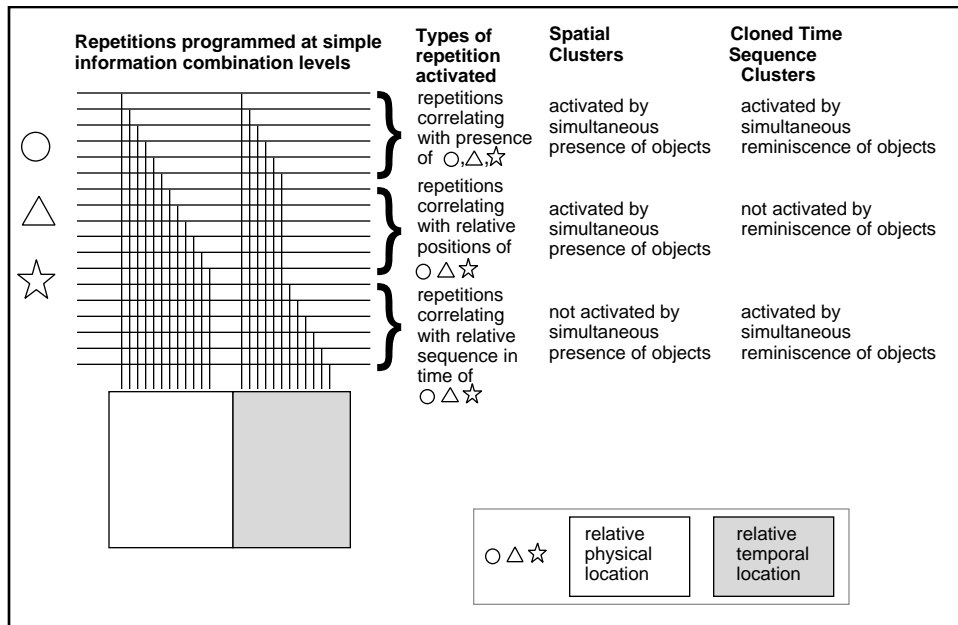


Figure 13 Cluster cloning to support time sequence dependent action recommendations. Clusters which detect the repetition of conditions in which objects are aligned in space have as input much of the information needed for clusters detecting time sequence repetitions. Additional information is required on relative time sequence in which repetitions were recorded. Such information could be derived from the resource map discussed in the text. If such information happened to be available at the space module, weak similarity could lead to imprinting, but because the new repetitions contained time information they would develop into a separate population cloned from the space population.

reminiscent activation of the perception of the three shapes occurred, then much of the input information which generally produced outputs from the cluster would be present, although not the relative position information. In the hippocampus model suggested in Coward 1990, information indicating relative time sequence would also be active. If some of this information happened to be connected to the input to the cluster by the biased random connection selection process, then imprinting could produce an output from the cluster which could become associated with appropriate time sequence related behavioral recommendations. Because the imprinted repetitions would include time sequence related information and not spatial information, two populations of imprinted repetitions would develop for spatial related and temporal related behavioral recommendations. The temporal cluster would have been cloned from the spatial cluster. This scenario suggests that there is a spatial basis for our thinking about time. Jaynes (1976) has emphasized the pervasiveness of spatial models and vocabulary in our thinking and speech about temporal relationships.

Given the heuristic process for associating clusters with behavior it is also possible to see how a useful cluster could be adopted in another function. For example, clusters which frequently generated appropriate behavior in the food seeking supercluster could be used as a template in the speech generation supercluster. A new cluster template would be established with the appropriate neurohormone sensitivities. The information which frequently generated outputs in the existing clusters would be connected as inputs to the new clusters. Outputs would be speech related behaviors and could include both experimental phoneme activation and pseudosensory input activation. Imprinting would be triggered in the new clusters whenever the existing clusters generated an output recommendation, and the outputs from the new clusters could be heuristically guided to appropriate behaviors.

8. The Recommendation Architecture In Biology

For the reasons discussed earlier, biological brains must experience strong selection pressures in favor of simple functional architectures. However, the arguments of appendix I demonstrate that there are only two options for such an architecture, and biological brains do not exhibit the memory, processing separations of systems with the instruction architecture and do not functionally resemble the instruction based operations of such systems. The evidence that biological brains have adopted the recommendation architecture is extensively discussed in Coward 1990 but can be summarized in three areas: functional evidence; structural evidence; and evidence from the functional effects of damage.

The functional evidence includes the ubiquitous role of categorization in animal (Lorenz 1977) and human (e.g. Harnad 1987) cognition. Further evidence is the existence of heuristic definition of functionality. The major requirement for management of the distribution of information in a heuristically defined functional architecture is consistent with the operations of dream sleep.

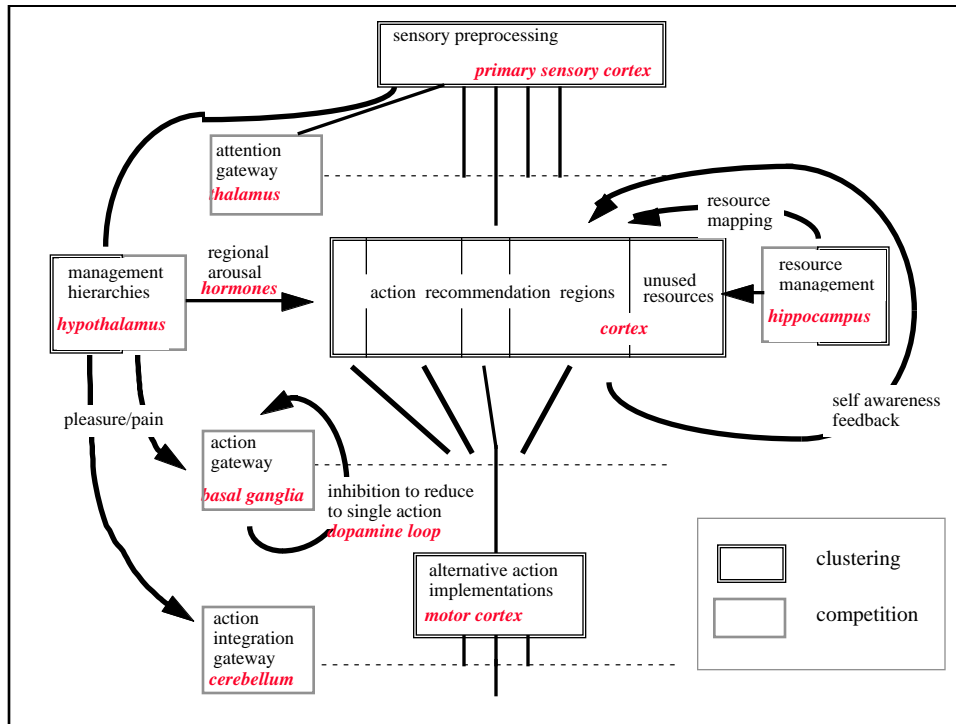


Figure 14 Major physiological structures in the brain can be mapped into similarity clustering and competition subsystems with an architecture resembling figure 3 plus subsystems as illustrated in figure 11. Information generated by an initial clustering in the primary sensory cortex is competitively reduced through the thalamus structures. The remaining information is reclustered in the cortex to generate alternative action recommendations. These alternatives complete in the basal ganglia structures, and the information making up one alternative is allowed to proceed from the cortex to the cerebellum where it is reclustered to specific muscle movements. Parallel similarity clustering subsystems manage major system management functions.

Physical evidence is based on the structures of a system with the recommendation architecture reflecting the constraints of that architecture, just as memory and processing structures appear in any instruction system. Figure 14 indicates, following Coward 1990, how major physiological structures in the brain can be understood as clustering and competition structures. At a more detailed level, cortex columns correspond with cluster modules (Coward 1990, Tanaka 1993). The strong layering observed in the cortex is as expected from the functional separations within clusters discussed earlier, with functional separations such as management of incoming information, indication of requirement for new cluster, indication of cluster membership, maintenance of repetition orthogonality, maintenance of sufficient diversity in cluster outputs to indicate different conditions within the cluster definition differently etc. leading to physical layering. De Felipe (1997) has described strong cross connectivity of spiny stellate cells in layer IV of the cortex, and inhibitive back projection by bouquet cells. This connectivity would be the type required to implement the cluster membership subfunctions discussed earlier (see figure 5), with the cross connectivity detecting repetition of a significant subset of the information combinations defining the cluster and stimulating imprinting of additional combinations, and the inhibitive backprojection preventing imprinting in the presence of already significant output. Another example is backward projections, for example from the higher to the lower areas of the primary sensory cortex which terminate in layer I (Cauller 1995) and which I have argued above are associated with maintaining repetition orthogonality under conditions of multiple repetitions per device. The observation that the typical cortex neuron has about 100 thousand inputs and fires if about 75 of them are active (Cauller 1997a) indicates that neurons are programmed with multiple repetitions. Similarity clustering by detection of information repetitions requires that primary connections are stimulative, although connections managing local recognition of similarity may be inhibitive. This requirement is consistent with observations that long range cortex connectivity is mainly stimulative, but a proportion of local connectivity is inhibitive (Douglas and Martin 1991).

In figure 14 there are a number of subsystems, also with the clustering, competition structure, which perform system management subfunctions. One of these subfunctions modulates the relative arousal of behavioral superclusters in response to perceived need. For example, within the subsystem clusters take inputs indicating internal body states and generate behavioral recommendations to lower device thresholds in a food seeking supercluster. Such subsystem clusters would roughly correspond with the cognitive category / recommendation *low blood sugar / feel hungry*. Clusters roughly corresponding with *threat to self or property / feel angry* could increase the probability of aggressive action recommendations. Clusters correlating with perception of attractive member of the opposite sex could generate sexual arousal recommendations increasing the probability of courting behavior. Following Coward (1990), the hypothalamus is identified as the structure which performs this function within the mammal brain. Another key subsystem clusters information about resource usage by cluster modules and generates recommendations to assign additional resources. Coward (1990) argued that this function is performed by the hippocampus using a map of cortex resources to assign additional resources as required. A side effect of such a mechanism is that the map provides an implicit time sequence of recorded repetitions.

The evidence from damage is particularly striking. In general, the types of functional deficits which can and cannot result from physical damage are a strong indication of system functional architecture. The heuristic similarity clustering process means that information extracted from a single experience will be recorded in numerous clusters located in a number of behavioral superclusters. Physically localized damage to the similarity clustering function is therefore unlikely to remove all such information, the effect of such damage will be reduced ability to generate recommendations of one or some types. This exactly corresponds with the observation that local damage to the cortex does not remove event memory, but may produce behavioral shifts (the classical reference is Harlow 1868). Damage affecting the management subsystem which assigns resources for recording repetitions may result in time sequence deficits if the damage interferes with access to a block of resources assigned over a period of time. Coward (1990) has pointed out that the deficits associated with Korsakov's syndrome, including loss of a time block of memory and the inability to create new memories, are exactly as expected from damage to the communication between a cortex performing the clustering function and a hippocampus performing the mapping function outlined above. A further example is damage to a competitive reduction function. Coward (1990) has argued that the symptoms of Tourette's syndrome, including rapid behavioral shifts, are consistent with a failure of the competitive system to reduce behavioral alternatives to one or less, and the dopamine function which is associated with Tourette's syndrome plausibly plays the role of damping down the number of alternatives in the basal ganglia.

The recommendation architecture supports a plausible evolutionary sequence by which cognition evolved. The scenario identifies a sequence of small physiological changes, each of which generates a change to system functionality which is of value to the system without any of the subsequent changes. In the scenario, the neuron emerged first within a multicellular organism with the ability to drive a response from combinations of sensory inputs including internal physiological conditions. Genetically programmed hierarchies of neurons developed to manage choices among complex motor response alternatives. Neurons extracting repetitions from internal physiology took on the role of tuning genetically hard wired neuron hierarchies to details of individual body development. This tuning was an early hierarchy management role and eventually evolved into pleasure and pain. Networks of randomly connected neurons were added which allowed repetitions extracted from individual experience to be used to generate a wider range of potential behavioral responses to a given situation. Pleasure and pain extended their functions to manage the selection probabilities of alternative behaviors. Dreaming put statistical bias on the random connectivity resulting in better use of resources. Speech allowed communication of internally generated clusters, and as a side effect made internal feedback possible, which allowed more extensive searching of individual memory to develop better behavioral responses in complex social situations. This feedback is experienced as self awareness (Jaynes 1976), and the search process is experienced as the explosion of associative images we call consciousness. More detail of the later cognitive developments is given in Coward 1997b and below.

The recommendation architecture paradigm thus demonstrates considerable explanatory power for the behavior of biological brains. A number of further experimental tests of the existence of the recommendation architecture in biological brains were proposed in Coward 1990. These tests apply to instances of the architecture which exhibit heuristic definition of functionality, i.e. mammal brains. One test is the existence of the imprinting mechanism as described earlier as a primary plasticity mechanism in mammal cortex neurons. A second test is the existence of the neuron connectivity changes resulting from

the information distribution management function of sleep. A third test is at the phenomenological level deriving from the information management function. The prediction is that under conditions of dream sleep deprivation accompanied by waking experience which is intensive with significant novel content in some behavioral domains where there is a biological need to take action, and minimal in other behavioral domains, then behavior in response to the intensive experience will shift towards types of behavior appropriate to domains which have not been stimulated. The reason is that imprinting resources are required to generate action, especially if some novelty is present. Dream sleep renews these resources where they have been depleted. If such resources are depleted in the most relevant clusters and a need to respond is present, less relevant cluster will be used resulting in less appropriate behavior.

9. High level functional definition of consciousness

The term consciousness is used to label phenomena ranging from the simple ability to respond to a stimulus to the ability to experience and talk about a constant succession of mental images generated largely independently of sensory input. These mental images can include objects never directly experienced. Images of self acting and experiencing can be generated, with either an internal or an external viewpoint (the "I" and "me" in Jaynes 1976). In Jaynes' example, we can generate a mental image of ourselves running to the lake shore and diving into the water, and also a mental image of the feel of our body swimming, the splashing of water on our face. The first of these images is generally not part of our direct sensory experience.

From a functional point of view, four levels of consciousness can be distinguished (Coward 1990), with each level including all the lower levels. At the lowest level 1 is the ability to respond to a stimulus, which does not require a nervous system. Indications of favorable or unfavorable conditions are simply stimuli, not a separate pleasure and pain system as is used at higher levels. Level 2 is the ability to develop complex, appropriate responses to complex combinations of stimuli. Level 2 requires a nervous system but does not require recording of individual specific experiences. A primitive separate pleasure and pain system may be used to tune the genetically programmed nervous system to variations in individual body growth. Level 2 is typical of insects and reptiles. Level 3 includes the ability to record the state of the brain during an individual experience and to partially reactivate that state in later, similar experiences as an aid to generation of behavior. Simultaneous activations of various combinations of individual specific memory traces correspond with alternative action recommendations. The choice between alternatives is managed by an extended pleasure and pain system which influences the probability of future action selections. Level 3 is typical of mammals. Level 4 adds to the ability to generate multiple behavioral alternatives the ability to generate from an alternative a set of pseudosensory inputs as if the alternative had been accepted and carried out. Additional alternatives are generated by the combination of actual and pseudosensory inputs, which can in turn result in further pseudosensory inputs and yet more behavioral alternatives. The process is experienced as a constant succession of mental images, and greatly extends the range of individual specific memory which can be searched for behavioral guidance in a given set of circumstances. Level 4 consciousness requires that a representation of a behavior be created and communicated to a point at which pseudosensory input corresponding with self performing the behavior can be generated (Coward 1990). This generation of a mental image by a symbolic representation strongly resembles speech, and both Jaynes (1976) and Dennett (1991) have argued that speech plays an important role in consciousness. Following Jaynes' phenomenological discussion, Coward (1990) suggested that physiological paths internalizing the route from utterance to hearing are the basis for level 4 consciousness. A process by which this internalization could have occurred is outlined in Coward 1997b.

From this point in the paper, the term consciousness without a level qualification will mean only level 4 consciousness.

Consciousness Models

Lorentz (1977) proposed that the phenomena of human consciousness derive from a number of capabilities which exist in other animals or have developed from such capabilities. The major change between other animals and human beings is in the way in which these capabilities interact. These capabilities are the abilities to extract constant objects from sensory input; to maintain mental models of relative positions; to treat every object as biologically relevant in order to find objects which are genuinely significant (i.e. curiosity); to compose complex actions from a vocabulary of muscle movements; to try past behaviors in new situations; and to imitate behavior. Humans have added the abilities to represent time

with spatial models and to imitate behavior by indirect instruction. Lorenz does not directly address the source of the constant stream of mental images, or discuss any neurophysiological mechanisms. The capabilities which Lorenz identifies are all typical of a recommendation architecture exhibiting heuristic definition of functionality. For example, the capability to extract constant objects from sensory input reflects the search for information combinations which repeat, which may require the inclusion of information about the state of sensory organs. Curiosity as defined by Lorenz reflects the need to associate similarity clusters with behaviors.

Jaynes (1976) offers a well developed psychological theory of consciousness along with some controversial ideas on the evolutionary and historical origins of consciousness. In Jaynes' view, as we interact with an object we perceive the object in many different ways over a succession of instants in time. These different ways Jaynes labels percepts. An object is represented in consciousness by a reactivation of a percept. A conscious image is thus the reproduction of one of the instantaneous mental images which occurred during direct perception of the object in the past. New objects are assimilated into existing categories developed to define behavioral responses. Mental objects are arranged in a spatial framework called a mindspace. This mental arrangement is used to represent relationships in space, in time, and in emotion (e.g. containing or letting out anger). Central to Jaynes' theory are the concepts of analog I and metaphor Me, which are conscious representations of self acting and self experiencing. The ability to tell stories is the basis for arranging events in a sequence and thus to modeling causal relationships. Consciousness for Jaynes is thus a linguistic invention, developed in response to a need for behavioral responses to extremely complex social situations. Implicit in Jaynes' theory is the recommendation architecture concept of development of repetition similarity clusters to define behavioral responses, and the behavioral role of consciousness in developing behavior in complex social situations developed in detail in this paper is the functional role originally proposed by Jaynes.

Baars (1988) has proposed a phenomenological theory of consciousness around the concept of the global workspace, and has begun (1994) to develop the corresponding physiological theory. According to Baars, the contents of consciousness are the contents of a global workspace located in the primary sensory projection areas of the cortex (e.g. VI for the visual sense). A set of input processors generate visual images, inner speech etc. which compete through structures like the brainstem reticular formation and the nucleus reticularis for access to the global workspace. The contents of the global workspace are distributed widely to specialty unconscious processors including perceptual analyzers, output systems, action systems, syntax systems, and planning and control systems.

In Baars' model, "the overall function of consciousness is to provide very widespread access to unconscious brain knowledge including autobiographical memory ... ; the lexicon of natural language ; automatic routines that control actions ; and, by way of sensory feedback, even the detailed firing of neurons and neuronal populations". The primary difference from the recommendation architecture approach is that there does not appear to be a clearly articulated view of the partitioning of behavioral functionality in Baars' model. Concepts like perceptual analyzers, output systems, action systems, planning and control systems are von Neumann concepts implying the use and generation of unambiguous system information. However, the input processors postulated by Baars have a functional resemblance to clusters in the sense that they generate alternative activation recommendations which compete for implementation.

Coward (1990) modified Jaynes' phenomenological theory of human consciousness in a number of areas and provided a physiological model of consciousness based on recommendation architecture theory of brain architecture. The first modification to Jaynes' theory is that direct perception is by extraction of a set of combinations of information from the perceived object, most of the combinations being repetitions of combinations created by perception of previous, similar objects, and a small set being created by the immediate process of perception. A mental image is the reactivation of a subset of the repetitions within a set of repetition similarity clusters frequently activated by similar objects. The second modification is that overlapping mental images can activate a set of pseudosensory repetitions which are frequently present during the perception of additional objects, which can in turn generate a mental image of the additional objects. These pseudosensory activations are developed and communicated using mechanisms similar to those used to generate speech, but are not themselves necessarily oral. The third modification is that perceptions and images of objects are functionally action recommendations, and multiple action recommendations for the same object result from the object being assigned to multiple sets of repetition similarity clusters which may be similar in sensory terms but separate functionally. The relative strength of action recommendations derived from the same object is modulated by feelings and emotions such as hunger, anger, and fear. Such feelings and emotions are themselves modulated by action recommendations.

The fourth modification is that selection of action is functionally separate from generation of action recommendations, and only the latter can enter consciousness. The 1990 model is developed in more functional detail in the body of this paper.

Metzinger (1996) has defined a set of components which need to be specified at both the psychological and physiological level in any integrated, systematic theory of the phenomenal content of consciousness. These include the smallest unit of phenomenal content; the smallest representable unit of content; object constitution, or binding; temporal relationships; spatial relationships; a hierarchy of relationships derived from spatial; self representation; first-person and other perspectives; and global modeling of reality which is not recognized as a model. The current paper attempts to define a complete set of components of consciousness and their interactions at both the psychological and physiological levels within a recommendation functional architecture and addresses all the functions called for by Metzinger.

10. Architecture of a System Delivering Consciousness

The highest level functional separations in a system delivering level 3 consciousness are shown in figure 3 (following Coward 1990). An initial similarity clustering extracts sensory system independent repetitions from raw sensory input. A competition between sets of repetitions from different sensory domains results in the repetitions from one domain proceeding to further clustering. This first competition is the attention process. The second clustering generates parallel activations corresponding with recommendations of different types of behavior (aggressive, fearful, sexual, food seeking, cooperative, etc.).

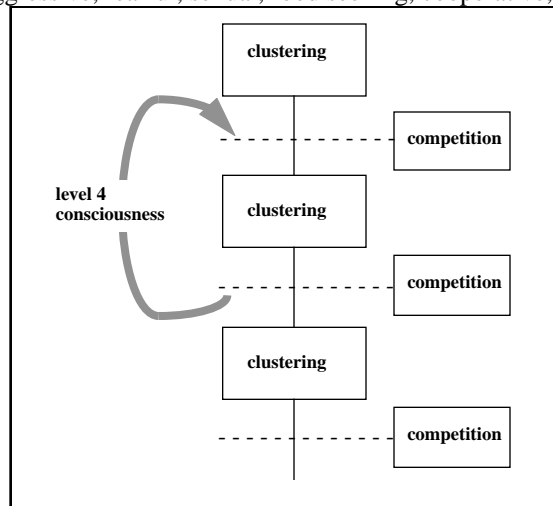


Figure 16 The functional feedback loop which makes possible the activation of mental images independent of direct sensory input

A competition between the different recommendations results in selection of a behavior. The activations corresponding with the selected behavior are further clustered into a portfolio of action sequences which compete to generate an integrated action.

The comment of Lorenz (1977) on the importance of the animal capability to extract constant objects from sensory input, and Jaynes' (1976) observation that most mammals assimilate slightly ambiguous perceived objects into previously learned schemas are both reflections of a functional architecture based on similarity. The importance of categorization in human cognition also reflects the recommendation architecture with sets of clusters operating through a competitive function to define cognitive categories.

The transition to level 4 consciousness depends upon the introduction of an additional functional mechanism, which is derived from signaling capabilities and illustrated in figure 16.




Figure 15 Signaling depends upon the signal being able to activate a significant subset of the local information which would be activated by the object signaled, and thus generate the same action recommendations

To demonstrate the development of this capability, consider how signaling achieves its behavioral effect in a recommendation architecture as illustrated in figure 15. The presence of an elephant results in a population of repetitions being generated which in turn generates a set of behavioral recommendations, including flight. If one of the behavioral recommendations can be to shout "elephant!" the value of the signal is its capability to generate flight or other appropriate behavior in the absence of direct sensory input. The simplest way to achieve this is if the sensory input of the word activates many of the intermediate level repetitions activated by direct perception. These pseudosensory repetitions are then clustered into the same behavioral recommendations as would result from direct perception. An appropriate point in level of complexity for these pseudosensory repetitions to be activated by signals would be more complex than raw visual repetitions but prior to the attention competition. Entry at that point would minimize the processing of information but allow independent decision on whether to pay attention to the signal.

The creation of the associations which activate the pseudosensory repetitions from the word can be a Hebbian type mechanism. Provided that the appropriate functional architectural framework is available, connections in the recipient brain for activation of the appropriate pseudosensory repetitions could be established heuristically by simultaneously seeing the object (i.e. activating the appropriate sensory level repetitions) and hearing the word (i.e. activating the appropriate auditory repetitions) and establishing connections between functional components which are frequently active at the same time.

However, there is a side effect of this mechanism. If the word "elephant" is spoken, it has the capability to activate pseudosensory repetitions in the mind of the speaker as if the elephant were being perceived, a feedback loop has been established. The simplest result of this side effect is that sensory images can be prolonged after the original stimulus has disappeared. A commonplace example is when we speak aloud to ourselves a telephone number we have just read, long enough to drive the motions of dialing. Coward (1997b) has argued that use of spoken signals in an analogous manner was important in making early human tool making less dependent on the copying of physical models (figure 17). The recommendation to speak the name of a tool could derive from associative activation from combinations imprinted when the tool had been used in the past, hence the recommendation to make a tool could be activated from conditions in which a tool was needed. The need conditions activate a mental image of the tool, which drives the tool making behaviors.




Figure 17 The set of active repetitions generated by looking at a tool are an important part of the information which generates the sequence of action recommendations to make a new tool. Once a word for a tool is available, the spoken name can generate the same set of active repetitions and generate the sequence of action recommendations to make a new tool in the absence of an example. The same mechanism makes it possible for an individual to accept oral commands. The words of the command generate the activated repetitions representing the commanded action, which can drive the action recommendations to carry out the action.

There is an important functional value gained if internal physiological routes replace the external spoken route, as illustrated in figure 18. When image activation is dependent on speech, the only images which can be evoked are those for which a word exists. In figure 18, a supercluster exists which generates recommendations to activate pseudosensory repetitions which have frequently occurred in the past when its component clusters have been producing outputs. Again, the connectivity to support the functionality could be established using a Hebbian type mechanism, and in general activation would be as a result of associative overlap. However, the internal function has the advantage that images of parts of the tool which have no name can be evoked and generate more detailed action recommendations. Much more detailed tools can be made, and furthermore the sustaining of images becomes easier, sometimes occurring without actual tool making. Because imprinting occurs when images are activated, whether or not actual tool making occurs, the additional repetitions will lead to more variability in the tools made. Greater variation in tools would therefore be expected from image driven tool making, and can explain the explosion in new types of tools observed in human archaeology in the 40,000 b.c. to 15,000 b.c. period (Jaynes 1976).

Such an internal feedback system can be the basis for a new function. As discussed earlier, the same sensory information will in general result in multiple alternative behavioral recommendations. For example, the result of attention being focused on a dog is the generation of recommendations of the friendly, aggressive, speech generation etc. types. If the feedback system is also applied to the friendly, aggressive etc. behavioral recommendations, the result will be activation of a set of pseudosensory repetitions which were frequently active in the past when such recommendations were generated, including when they were accepted, and could include repetitions imprinted by past internal body information as well as past external environment information. Such activations could therefore be experienced as self actions. Clusters developed from observation of other individuals could be cloned if such cloned clusters proved valuable in directing activation feedback. The activation of such cloned clusters would be experienced by the system as observation of self taking an action, viewed from outside of self. The result would be to activate an additional set of clusters and therefore generate additional behavioral

recommendations, and the functional role is therefore to search a more extensive range of individual specific memories for behavioral guidance.

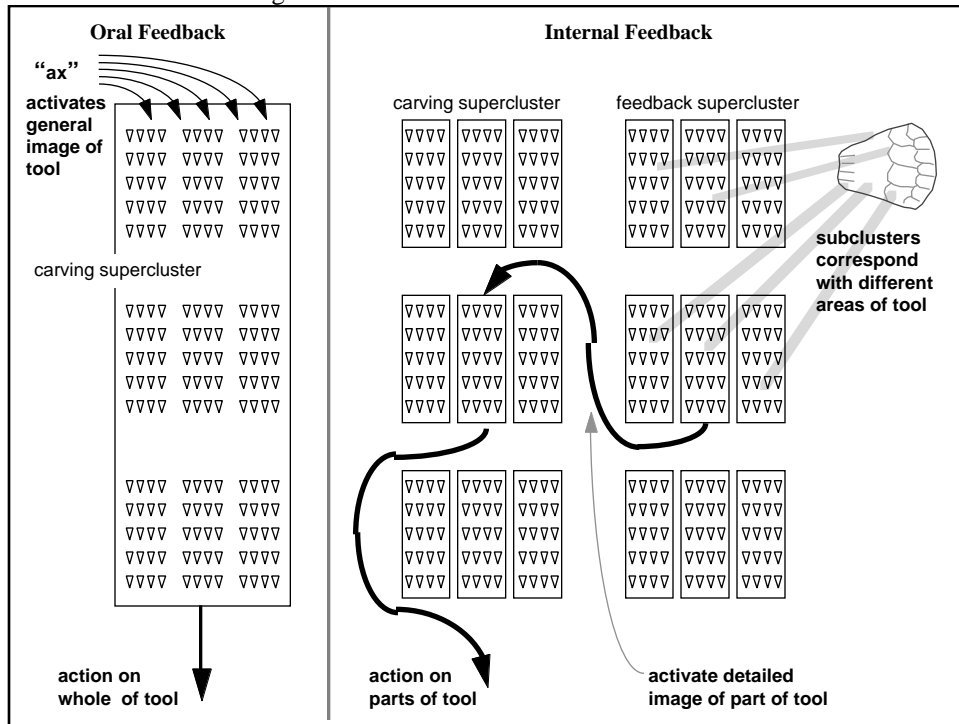


Figure 18 Comparison of orally controlled and internally controlled feedback. Words can only activate images for which there is a word. Internal feedback loops can activate subclusters for which there are no verbal labels but which can generate more detailed action recommendations. The internal feedback clusters are initially activated by associative overlap.

As Jaynes (1976) has argued, the primary value for this process lies in generating behavior appropriate in complex social situations as illustrated in figure 19. In that figure the brain has developed clusters which generate action recommendations when experiencing one particular individual in different circumstances. When this individual is experienced in novel circumstances, the first recommendation generated is one appropriate to the most similar of the programmed circumstances. This recommendation results in activation of pseudosensory inputs as if the action had been taken, generating enough repetitions to activate a recommendation appropriate to a second programmed circumstance. The overlap of the recommendations results in a somewhat different accepted recommendation, and the imprinted repetitions form the starting point for a set of clusters appropriate to the new circumstances. The feedback driven activation process results in repetition imprinting or even new cluster creation just as in activation by direct experience, and these virtual memories can also become resources for future behavioral guidance.

The feedback functionality is derived from speech and can therefore be expected to be unique to human beings, and the physiological location of the functionality would also be expected to be unique. Two possibilities are the inferior parietal lobe (Coward 1990) or the anterior cingulate cortex (a different interpretation of results discussed in Cotterill 1995).

When the feedback capability is combined with the associatively generated activations discussed earlier, weak associative activations can be recommendations to generate the full activations of the clusters usually activated by the associatively activated object. Such full activations can themselves give rise to further weak activations amplified in turn by feedback, resulting in the system experiencing a series of states similar to those resulting from a series of actual perceptions. The system thus experiences a series of images decoupled from direct sensory input.

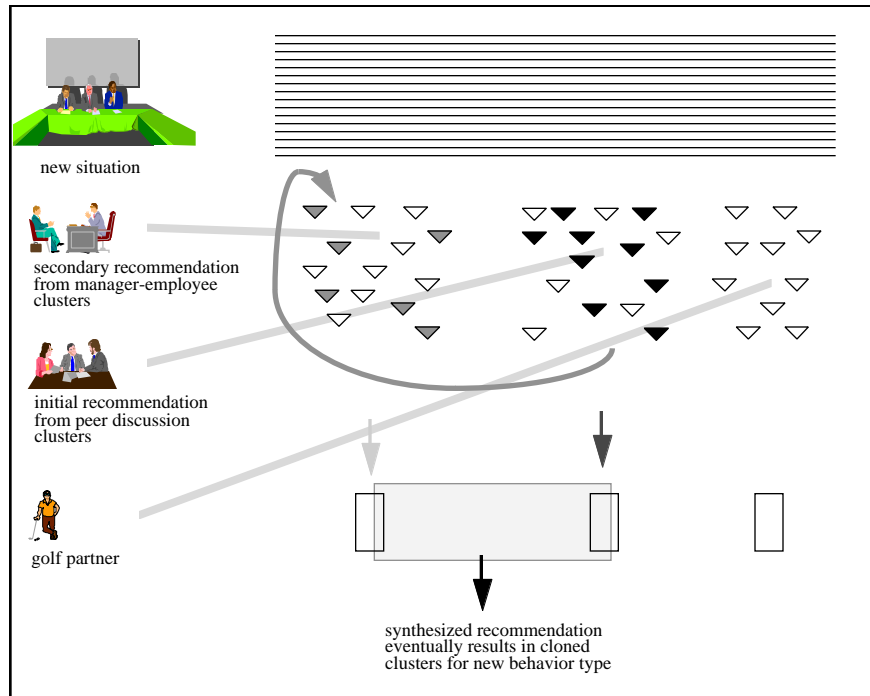


Figure 19 Initial action recommendations are derived from the most similar of previously defined cluster sets. Activation of simple repetitions frequently active when the higher order repetitions corresponding with the recommendation have been active in the past activates a wider range of recommendations. An action recommendation representing a synthesis results. Imprinting in all active cluster sets results in cloning a new cluster set for the new behavioral type.

It has been implicitly assumed that repetitions making up mental images of different objects can overlap to generate secondary repetitions which in turn by feedback generate images further objects as illustrated in figure 20. An important issue is the management of repetitions extracted from different perceptual objects. The functional need is to ensure that appropriate behavioral recommendations can be generated with respect to individual objects, but also that it be possible for activation sets for multiple objects to overlap and generate activation sets for additional objects, which must themselves be distinguished from their sources. This is essentially the binding problem and a widely discussed solution uses coherent modulation of firing rate across populations of activations (see Llinas et alii 1994). A more general solution is that the repetitions associated with a particular object are tagged with a particular firing modulation frequency in the 40 Hz range, the primary object of attention is identified by the main frequency, and other objects can be identified by modulation frequencies offset from the main frequency. The experience of the images of consciousness can then be understood as follows. There is a competition as described in Coward (1997a) between the ambiguous equivalent of the primal sketches proposed by Marr (1982) from which the winner gains access for all the repetitions extracted within the domain defined by the sketch into the behavioral recommendation generation function. All repetitions derived from the winning domain are tagged by the main modulation frequency, and it is output repetitions with that frequency which form the externally directed behavioral recommendations. Secondary objects may gain some access, but repetitions extracted from within their domains are tagged with a frequency offset from the main frequency. Output activation sets composed of offset frequency repetitions can only be recommendations to activate pseudosensory activations. Such recommendations, if accepted, will generate activation sets tagged with a single modulation frequency. This frequency could in some cases be the main frequency, as when an imagined object drives behavior, for example when writing a paper.

A critical capability is therefore the number of independent modulation frequencies which can be supported in the alternative behavioral generation function. The number is likely to be small, and defines the number of independent objects which can be imagined at the same time. The capability is only relevant for sustaining the feedback activation of images, and can therefore be expected to be limited to the human beings among Earth species.




Figure 20 Local information extracted from objects A1 and C1 is introduced simultaneously into a similarity clustering function. Confusion between the objects is avoided because their information is tagged by different rates of change of neuron firing rate, and recommendations generated are similarly tagged. The simultaneous presence of repetitions from the two objects activates a set of higher complexity repetitions extracted in the past from objects in cluster sets B, which have frequently been present at the same time as objects in cluster sets A and C. This set can be an action recommendation to activate the simpler repetitions which have frequently been present at the same time. Acceptance of the recommendation would activate an image of the object B1.

11. Options for Similarity Cluster Functionality

At this point it is useful to revisit the range of possible functions which a similarity cluster can fulfill. A cluster is any set of repetitions which can have a useful internal or external function, with the proviso that if the distribution of information between clusters becomes too complex, the system will become difficult to construct and maintain and the organism possessing such a brain will probably become extinct.

The functional role of a cluster is to generate behavioral recommendations. However, recommendations can also be directed at the system itself. A cluster composed of an appropriate set of repetitions can drive necessary system functions. Useful clusters could include detection that a particular activation set is driven by internal feedback and does not correspond with any current external object. Such clusters could recommend against certain types of behavior. Another type of cluster could detect that recommendations corresponding with similar sets of cluster activations sometimes produced favorable results, sometimes unfavorable. Such clusters could recommend reclustering of input information to find a consistent partitioning. An example could be if the same behavior towards a particular individual produced different results at different times, and activation of 'contradiction detection' clusters resulted in adding information such as time of day or mood indications to the inputs utilized and repeating a heuristic similarity clustering process.

There is clear behavioral advantage is such a 'presence of contradiction' clusters in terms of functionality in developing better strategies in complex social situations. Physiological mutations in favor of such a functionality would therefore be selected in evolutionary terms. However, once such a functionality exists in combination with the ability to sustain multiple independent mental images, the functionality which will tend to generate consistent belief systems such as history, philosophy, and science is in place. To understand why this is the case, consider the scenarios illustrated in figure 21.

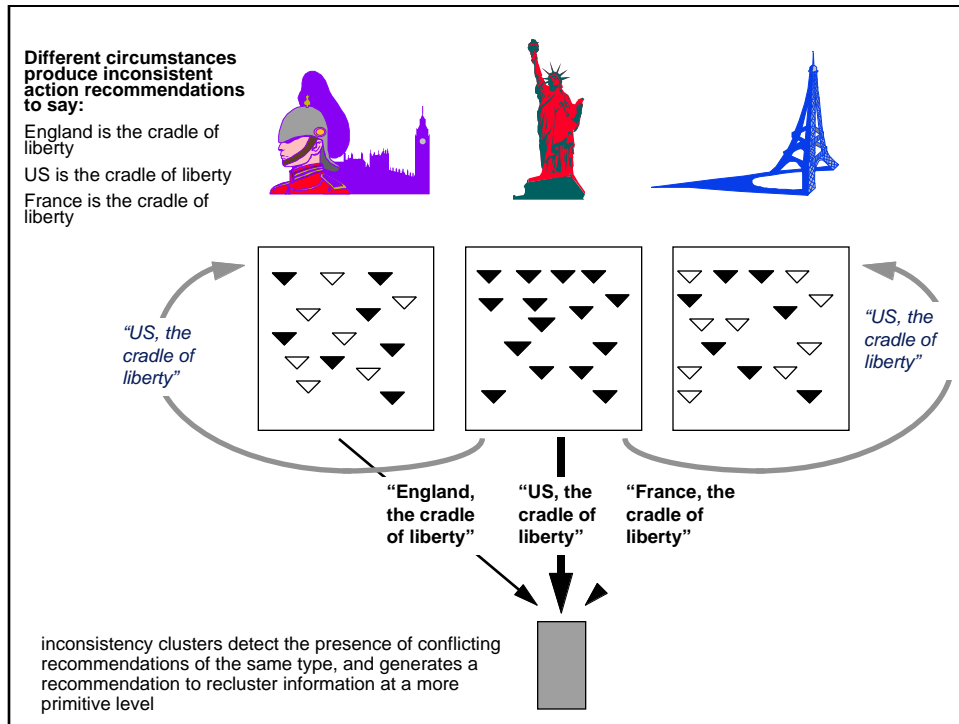


Figure 21 Without feedback, a brain can generate recommendations in different situations which would be contradictory if generated at the same time. Feedback results in the recommendation to say “the US is the cradle of liberty” in one situation activating enough simpler information to activate “England is the cradle of liberty” and “France is the cradle of liberty” recommendations at the same time. The simultaneous generation of conflicting recommendations of the same type is a condition which activates clusters with output recommendations to recluster the information at the simple level.

A human visits London, Paris, and New York. In London he sees the Houses of Parliament, and the sensory input generates the behavioral recommendation to say "London, the birthplace of liberty". The same individual visits Paris, sees the sites of the Revolution, and again the sensory input generates a behavioral recommendation, in this case to say "Paris, the birthplace of liberty". Finally, the same individual visits New York and sees the Statue of Liberty, and yet again the sensory input generates a behavioral recommendation, in this case to say "America, the birthplace of liberty". Note the similarity with early classical Greeks who could place the birthplace of Apollo at the sites of a number of different cities without apparent awareness of a problem (Johnson 1987). Now suppose that associative activation of secondary images, the ability to sustain multiple images, and inconsistency clusters are in place. When the individual says or thinks "New York, the birthplace of liberty" the 'birthplace of liberty' content may activate repetitions imprinted at the times when he has said the same things about the other places. These activations simultaneously generate recommendations to say the three contradictory statements. Detection of the contradiction generates a recommendation to recluster the input information towards a set of clusters which generate consistent statements, the first step towards a consistent historical view.

Other clusters could indicate that it was becoming more difficult to generate behavioral recommendations in some behavioral clusters, with action recommendations to add imprinting resources or modify similarity parameters etc. Yet other clusters could indicate the relative time at which resources used to imprint different images were assigned, thus giving the capability to generate behavioral recommendations based on the relative time of different experiences. Coward (1990) has suggested that the hippocampus performs the function of assigning cortex resources, and argued that the nature of the Korsakov's Syndrome deficit provided evidence in support of that view. The hippocampus could then be an important source of information on relative time sequence.

12. The Dynamic Experience of Consciousness

At any instant in time, sensory input is directed to two parallel similarity clustering functions. One function extracts repetitions which indicate the presence of coherent objects in the external environment ("primal sketches"). The other extracts repetitions indicating the presence of object characteristics (e.g. object color independent of illumination). Outputs from the first function are action recommendations to *pay detailed attention to this sensory domain* or *carry out programmed action* with respect to the primal sketch. Programmed actions include rapid response behaviors such as ducking, grabbing, etc. A *pay detailed attention* recommendation, if accepted, results in repetitions extracted by the second function from sensory input within the domain defined by the primal sketch gaining access to a further similarity clustering which generates alternative action recommendations with respect to the primal sketch. The tagging of such repetitions with a single modulation frequency overlaid on the neuron firing rates which indicate the presence of the repetitions is the mechanism which indicates to later functions that all the repetitions derive from one primal sketch. The mechanisms of this paragraph underlie the psychological experience of attention.

The alternative action clustering function is subdivided into superclusters which generate behavioral recommendations of different types. Each supercluster contains a hierarchy of clusters. Types of behavior include aggressiv, fearful and food-seeking for example. The relative probability of action recommendations of different types is modulated by a separate clustering function which clusters a combination of inputs including a significant proportion of inputs derived from within the body. For example, repetitions indicating body strength and external threat increase the probability of aggressive behavior, inputs indicating low blood sugar increase the probability of food seeking behavior. This modulation function underlies the experiences of emotion. Within one supercluster, recommendations can include shifting the focus of attention, for example to subdomains withing the current primal sketch (e.g. pay attention to the mouth and teeth areas of a dog). Recommendations could also include to increase the arousal of the supercluster itself (e.g. to become more angry).

One supercluster within the alternative action clustering function can generate speech recommendations, including in the simplest case to speak the name of the primal sketch. Note that as discussed earlier, a set of ambiguous clusters acting through a competitive function form the operational definition of a cognitive category. Another supercluster generates recommendations to activate inputs which were frequently present in the past when the current outputs were present. Such recommendations generated by pseudosensory input compete with *pay detailed attention* recommendations generated from current sensory input, and if accepted result in an activation state similar to the state resulting from sensory conditions in the past which contributed to imprinting of the currently active outputs. This feedback process underlies the experience of mental images independent of current sensory input. Such images can be distinguished from activations created by current sensory input both by different modulation frequencies and by the absence of the repetitions closest to sensory input.

Feedback from clusters imprinted by conditions including past conditions of the body generates mental images of self separate from actual current sensory input from the body, and may include images which could not be generated from any actual input (e.g. self viewed from outside). Mentalimages of self can include self conditions such as anger or fear.

As an active assembly decays, a new assembly can be activated by an action recommendation derived either from a primal sketch or from feedback of an associative activation. The new assembly of activated repetitions is perceived as a new single entity because all repetitions are tagged with the same firing rate modulation frequency, and generates a new set of action recommendations. In addition, some repetitions composed of repetitions activated in both new and old assemblies may be stimulated enough by the successive, perhaps multiple, assemblies to be activated. Such secondary activations can also generate action recommendations. Once an assembly has fully died away, it can only be reactivated by appropriate sensory input or appropriate pseudosensory input generated by assembly overlap. Some assemblies may not be accessible in a given situation.

What are the content limitations of consciousness? To answer this question it is important to appreciate both how the contents of consciousness are activated and how those contents can be determined by an outside observer. Baars (1996) offers a practical definition of human consciousness as including ".....the ability to answer questions, report perceptual events". This definition is consistent with level 4 consciousness but only addresses how the contents can be determined. Consider how split brain patients can be shown (right hemisphere only) an embarrassing picture, and respond with embarrassment while being unable to explain the actual cause, often inventing reasons to account for their reaction. The physical breakdown is that information about the picture has not been able to reach the speech centers of

the left hemisphere, and we typically conclude that there is no consciousness of the picture, although repetitions derived from the picture must be active in the right hemisphere and split brain patients exhibit apparently normal consciousness in less constrained conditions. Any mammal with level 3 consciousness has a set of repetitions activated by a sensory experience which are a mental image which can even generate communication recommendations in some cases. The distinctive property of level 4 consciousness is the ability to generate action recommendations to create pseudosensory inputs as if a non-perceived condition were present. The brain function which can generate such recommendations is derived from the function generating speech recommendations and both functions are located within the clustering function which generates all action recommendations. An accepted image generation recommendation produces pseudosensory inputs to that same clustering function, which could in turn generate speech recommendations to communicate current content. Image generation recommendations cannot directly generate pseudosensory inputs to any competitive function or to the initial sensory preprocessing clustering function, and these functions cannot generate speech recommendations. The contents of these functions are therefore outside of the scope of consciousness.

Consider the state of the brain when paying attention to a coffee mug. An assembly of parallel cascades of neuron firing extends through the cortex from visual input to action recommendations such as 'pick up and drink'. Imagining a mug uses a recommendation generated internally by assembly overlap to activate by feedback an assembly of neuron firing which extends for part of the length of the assembly generated by viewing a real mug. If the point at which the feedback initiates the assembly is close to the senses, the image generated is strongly visual. If further away the image is more abstract. If one of the parallel cascades generates a speech action recommendation, we can verbalize the content of the assembly at that point, such as an action recommendation to 'pick up the mug'.

Consider now the process of molding a mug out of clay. An assembly generated by paying attention to the clay is followed in attention by an assembly made up of repetitions extracted from past mug experiences which is generated by feedback of an associative activation. Repetitions activated by overlap of the two assemblies generate action recommendations to manipulate the clay. Other associatively generated activations are action recommendations to focus attention on detailed parts of the mug, say the handle, and so generate action recommendations which affect the appropriate area of the clay. Activation of unrelated assemblies during the molding process or when thinking about it may influence the repetitions activated and thus the result. Imprinting of additional information combinations in the course of generating action recommendations makes it possible to regenerate the action recommendation in the future and experience reminiscence.

Suppose now that assembly overlap, associative activation, recommendation acceptance and feedback generates a set of activations characteristic of a chimpanzee soon after a mug assembly. Repetitions activated by the overlap of the succession of mug and chimpanzee assemblies generate the associative activation of a relationship: chimpanzee throwing mug. Feedback generates an assembly corresponding with actual observation of the action. If similarity clusters corresponding with "I" were activated close to this time, overlap could generate pseudosensory input of self throwing the mug, along with an action recommendation to do it. Within a ball-playing supercluster this could be a strong recommendation. However, feedback of "I" taking the action generates an assembly which inputs to clusters constructed from social experience. Alternative action recommendations of 'put the mug down carefully' are generated along with 'inhibit throwing the mug' from clusters extracting inconsistencies. In addition, generation of pseudosensory representations of "me" experience during an action of this type could be fed back to generate the repetitions associated with attention on my internal experience associated with such an action, generating additional positive and negative action recommendations. In complex social situations, this process of extraction and feedback may continue through many cycles of attention, supplemented by action recommendations to 'become angry' etc. which change the type of action recommendations generated. Clusters recording a wide range of experience are searched, with additional repetitions imprinted, until an action recommendation is accepted and implemented.

There is no qualitative difference between direct sensory experience and reminiscence and stream of consciousness, all are constructed from the activation of widely shared repetitions. However, direct experience activates repetitions close to sensory input which are generally not activated in imagined images. Clusters indicating the absence of such repetitions could generate recommendations to behave accordingly, enabling a behavioral distinction between direct experience and reminiscent experience. In dream sleep the interaction of assemblies to generate action recommendations, and thus create a permanent record, is non-normal, as demonstrated by the relative rarity of remembered dream compared

with total dream. Dream sleep has the function of limiting the distribution of information between functional components at every level but making it available where it has the highest probability of being functionally relevant. At the device level this role consists of programming additional information combinations to be capable of imprinting the appropriate subsets to enable activations in response to new objects and experiences in subsequent wake periods. Actual imprinting is therefore turned off during this preparation period. A fast rerun of past repetition activations is used to identify combinations of information frequently activated at the same time, which become potential additional repetitions or distribution routes. The anomalous nature of the impressions created when imprinted (and therefore rememberable) sets of repetitions are created in dreaming is because the acceleration results in combinations of activated information which would not occur in regular waking experience, but recorded if mid-dream waking turns on imprinting.

13. The Functional Components of Consciousness

As discussed earlier, Metzinger (1996) has defined a set of components which need to be specified at both the psychological and physiological level in any integrated, systematic theory of the phenomenal content of consciousness. These include the smallest unit of phenomenal content; the smallest representable unit of content; object constitution, or binding; temporal relationships; spatial relationships; a hierarchy of relationships derived from spatial; self representation; first-person and other perspectives; and global modeling of reality which is not recognized as a model. In terms of the recommendation architecture model described in this paper, the smallest units of phenomenal content are the imprinted repetitions at the device level which can be activated by feedback of associative activations. The smallest representable units of content are sets of clusters which can generate recommendations to activate the population of repetitions which most often are present when the cluster set is activated. Object binding is indicated by a common modulation of firing rate shared across a population of activated device level repetitions. Any temporal or spatial relationship is a cluster or set of clusters with an associated type of behavioral recommendation. Self representation occurs when a set of clusters created by imprinting of information combinations occurring in past experience of self are activated, plus clusters cloned from experiences of other individuals. Activation of these clusters has the functional role of generating pseudosensory information as if a behavioral recommendation had been implemented, in order to search a more extensive range of individual specific memory for further behavioral recommendations. Global modeling of reality which is not recognized as a model corresponds with the hierarchy of clusters heuristically created in the course of experience which search for repetitions in current experience to generate behavioral recommendations. The hierarchies of clusters are an implicit model of reality, but activations within the "model" are driven by information derived from the current environment and directly generate recommendations for action on that environment. The "model" is not a parallel path which can be compared with reality and thus is not experienced as a model.

The functional theory as described is radically different from content focused models such as the global workplace theory proposed by Baars (1988), because Baars' model incorporates paradigms based upon the use of unambiguous information. In Baars' model the output from a wide range of specialized processors competes for access to a small global workspace, and the contents of that workspace are widely distributed throughout the brain. Only a limited number of specialist processors can broadcast global messages via the workspace because different messages may be contradictory. The contents of consciousness are the contents of the global workspace and become conscious by being broadcast. In the theory outlined here, there is no small "executive" workspace or working memory. The contents of consciousness are the currently active information repetitions wherever they are physically located in the alternative behavioral recommendation function (which includes a large proportion of the cortex). There is competition between different sets of information for access to the behavioral generation function, but the winner of this competition is not the content of consciousness, it drives the activation of more complex repetitions which are the contents. Because multiple parallel behavioral recommendations are generated in response to any winning information combination, and in general these parallel recommendations include recommendations to speak, it is possible to talk about the current configuration of activations. Some clusters are in a sense specialist processors which compete for access, but there are two major types. One type is primal sketch clusters which like Baars' processors are outside of consciousness. The other type is clusters corresponding with recommendations to activate information in the alternative behavioral generation function. These clusters are within the potential contents of consciousness. As an example on the

psychological level, I can sometimes be aware of and speak about an inclination to activate an image before or at the same time as activating it.

Baars' model places emphasis on internal consistency as a key restriction in determining which combinations of specialized processors gain access, but does not address the mechanisms which would enable consistency to be achieved. In the model described here, consistency is imposed because the only repetitions which will be activated by the input information are ones which have frequently occurred together in the past, and feedback loops cause the system to converge on a self consistent set.

Baars' model does not address the issue of how qualia arise from the information processing in the brain. The explanation in the proposed model can be understood by considering how the experience of the color red develops. Repeated experiences of objects with the color red in different situations will result in heuristic definition of clusters. The heuristic process searches for information repetitions. The profile of light wavelengths reaching the retina will not repeat even for the same object because of differences in ambient light. Repetitions will occur if combinations of information from both ambient light and object light are included, and clusters will develop on this basis. The experience of red corresponds with these clusters being activated. There could of course be some genetically controlled bias in initial cluster connectivity which favors development of useful clusters, particularly in the sensory preprocessing clustering function, but because the clusters are created at least partially heuristically, they will incorporate information extracted from experiences which included the color red, and the activation of such information can lead to activation of other clusters activated in such experiences by the feedback process. Such activations can therefore in some cases become part of the experience of the color red.

14. Conclusions

There are only two possible qualitatively different types of simple functional architecture, the instruction architecture based on consistent use of instruction functional elements exchanging unambiguous information, and the recommendation architecture based on consistent use of repetition/action recommendation functional elements exchanging ambiguous information. Systems with the recommendation architecture can have the ability to heuristically define their own functionality. A theory has been described that biological brains have been constrained by selection pressures to adopt simple functional architectures of the recommendation type, and the hierarchy of functional separations which exists as a result can be the basis for understanding cognitive phenomena in terms of physiology. The existence of a simple recommendation based functional architecture means that the human experience of a constant succession of mental images, including self images, independent of sensory input can be understood at high level as the experience of a function which activates potentially relevant information recorded from an extensive range of individual experience for behavioral guidance in complex situations. Such a function can activate a wider range of information than the information activated directly by sensory input. The operation of this function can be traced through a hierarchy of functional separations at increasing levels of detail down to the operations of individual neurons. To be effective in understanding cognitive functions, functional modules must be components in a recommendation architecture. Physiological structures can be understood in terms of major functional separations, and patterns of neural activation such as sleep with dreaming understood in terms of the execution of recommendation architecture functions. The relationship between physical damage and defects and the resulting system deficits can be understood through the functional hierarchy. Philosophical theories can be evaluated for consistency with the underlying biological architecture. Electronic systems with phenomenology very similar to human psychology could be designed and built using the recommendation architecture paradigm.

References

- Baars, B.J. (1988), *A Cognitive Theory of Consciousness* (Cambridge: Cambridge University Press).
- Baars, B.J. (1994), 'A Neurobiological Interpretation of Global Workspace Theory' in *'Consciousness in Philosophy and Cognitive Neuroscience'* ed. Revonsuo, A. and Kamppinen, M. (New Jersey: Erlbaum).
- Baars, B.J. (1996), Understanding Subjectivity: Global Workspace Theory and the Resurrection of the Observing Self, *Journal of Consciousness Studies*, 3, (3)
- Bressler, S. L. (1994) Dynamic self-organization in the brain as observed by transient cortical coherence, in: Pribam, K. (ed.) *Origins: brain and self-organization*, Hillsdale, NJ: Erlbaum.
- Carpenter, G.A. and Grossberg, S. (1988), The ART of Adaptive Pattern Recognition by a Self-Organizing Neural Network, *IEEE Computer*, 3, 77-88

- Cauler, L. (1995), Layer I of primary sensory neocortex: Where top-down converges with bottom-up, *Behavioural Brain Research* 71: 163-170
- Cauler, L. (1997) *NeuroInteractivism: Explaining Emergence without Representation*, to be published.
- Cauler, L. (1997a) Private communication.
- Cotterill, R. (1995), 'On the unity of conscious experience', *Journal of Consciousness Studies*, 2, (4), pp 290-311.
- Cottrell, G.W. and Metcalfe, J. (1991), Empath: Face, emotion and gender recognition using holons, in *Advances in Neural Information Processing Systems 3*, eds. Lippmann, R., Moody, J., and Touretzky, D. S., San Mateo, 557-564, CA: Morgan Kaufmann
- Coward, L. A. (1990) *Pattern Thinking*, New York: Praeger.
- Coward, L.A. (1996) *Understanding of Consciousness through Application of Techniques for Design of Extremely Complex Electronic Systems*, presented at *Towards a Science of Consciousness*, Tucson, Arizona.
- Coward, L. A. (1997a) *The Pattern Extraction Hierarchy Architecture: a Connectionist Alternative to the von Neumann Architecture*, in *Biological and Artificial Computation: from Neuroscience to Technology*, eds. Mira., J., Morenzo-Diaz, R., and Cabestanz, J., 634-43, Berlin: Springer.
- Coward, L.A. (1997b) *Unguided Categorization, Direct and Symbolic Representation, and Evolution of Cognition in a Modified Connectionist Theory*, in *Proceedings of the International Conference New Trends in Cognitive Science*, eds. Riegler, A. and Peschl, M., 139-146, Vienna: Austrian Society of Cognitive Science.
- Coward, L.A. (1998) *A Functional Architecture Approach to Neural Systems*, to be published.
- De Felipe, J., (1997) *Microcircuits in the Brain*, in *Biological and Artificial Computation: from Neuroscience to Technology*, eds. Mira., J., Morenzo-Diaz, R., and Cabestanz, J., 195-206, Springer: Berlin.
- Dennett, D.C. (1991), *Consciousness Explained'* (Boston: Little John).
- Douglas, R.J. and Martin, K.A. (1991), 'A Functional Microcircuit for Cat Visual Cortex', *Journal of Physiology*, 440, 735-769
- Harlow, T. M. (1868) 'Recovery from passage of an iron bar through the head', *New England Medical Society* 2, 327-46
- Harnad, S. (1987), *Category Induction and Representation*, in *Categorical Perception: The Groundwork of Cognition* (ed. Harnad, S), New York: Cambridge University Press.
- Hebb, D. C. (1949) *The Organization of Behaviour*, New York: Wiley
- Jaynes, J (1976), *The Origin of Consciousness in the Breakdown of the Bicameral Mind'* (Boston: Harvard).
- Johnson, D.M. (1987), *The Greek Origins of Belief*, *American Philosophical Quarterly* 24, 4, 319-27.
- Llinas, R., Ribary, U., Joliot, M., and Wang, X.-J. (1994), 'Content and Context in Temporal Thalamocortical Binding' in *Temporal Coding in the Brain* ed. Buzsaki G. et alii., Berlin: Springer.
- Lorenz, K. (1977), *Behind the Mirror*, (New York: Methuen).
- Marr, D. (1982), *Vision* , New York: W.H. Freeman.
- Mendel, J. M. (1995), *Fuzzy Logic Systems and Qualitative Knowledge*, in *The Handbook of Brain Theory and Neural Networks*, ed. Arbib, M, Boston: MIT Press
- Metzinger, T. (1996), *Towards the Neural and Functional Correlates of Phenomenal Content ASSC Electronic Seminar*
- Mira, J., Herrero, J.C., and Delgado A.E. (1997) *A Generic Formulation of Neural Nets as a Model of Parallel and Self Programming Computation*, in *Biological and Artificial Computation: from Neuroscience to Technology*, eds. Mira., J., Morenzo-Diaz, R., and Cabestanz, J., 195-206, Springer: Berlin.
- Newell, A. and Simon, H.A. (1976) *Human Problem Solving*, Englewood Cliffs N.J.: Prentice Hall
- Pfeifer (1996), *Symbols, Patterns and Behavior: Beyond the Information Processing Metaphor*, in *Encyclopedia of Microcomputers*, 17, 253-75, New York: Marcel Dekker.
- Ritter, H. (1995) *Self-Organizing Feature Maps: Kohonen Maps*, in *The Handbook of Brain Theory and Neural Networks*, ed. Arbib, M, Boston: MIT Press
- Rosenblatt, F. (1961) *'Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms'* Washington D.C.: Spartan
- Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986), *Learning internal representations by error propagation*", in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Volume 1, eds. Rumelhart, D.E. and McClelland, J. L., 318-362, Cambridge, MA: The MIT Press.
- Sun, R. (1995), *Robust Reasoning: integrating rule-based and similarity based reasoning*, *Artificial Intelligence*, 75, 2, 241-296.
- Sun, R. and Peterson, T. (1997), *A Hybrid Model for Learning Sequential Navigation*, *Proceedings of IEEE CIRA*
- Tanaka, K., (1993), *Neuronal Mechanisms of Object Recognition*, *Science*, 262, 685-88.
- Taylor, J.G. and Alavi, F.N., (1993), 'Mathematical Analysis of a Competitive Network for Attention' in *Mathematical Approaches to Neural Networks* ed. J.G. Taylor (Elsevier)..