A Computational Model of the Human Visual Cortex

James S. Albus*

Senior NIST Fellow
Intelligent Systems Division
National Institute of Standards and Technology
1 Bureau Drive
Gaithersburg, Maryland 20899

ABSTRACT

The brain is first and foremost a control system that is capable of building an internal representation of the external world, and using this representation to make decisions, set goals and priorities, formulate plans, and control behavior with intent to achieve its goals.

The computational model proposed here assumes that this internal representation resides in arrays of cortical columns. More specifically, it models each cortical hypercolumn together with its underlying thalamic nuclei as a Fundamental Computational Unit (FCU) consisting of a frame-like data structure (containing attributes and pointers) plus the computational processes and mechanisms required to maintain it.

In sensory-processing areas of the brain, FCUs enable segmentation, grouping, and classification. Pointers stored in FCU frames link pixels and signals to objects and events in situations and episodes that are overlaid with meaning and emotional values. In behavior-generating areas of the brain, FCUs make decisions, set goals and priorities, generate plans, and control behavior. Pointers are used to define rules, grammars, procedures, plans, and behaviors.

It is suggested that it may be possible to reverse engineer the human brain at the FCU level of fidelity using next-generation massively parallel computer hardware and software.

Key Words: computational modeling, human cortex, brain modeling, reverse engineering the brain, image processing, perception, segmentation, knowledge representation

1. INTRODUCTION

Much is now known and progress is rapid in the neuroscience and brain modeling fields regarding how the brain functions.[2, 5, 7, 8] Much is known in the computer science and intelligent systems engineering fields about how to embed knowledge in computer systems.[1] Researchers in robotics, automation, and control systems have learned how to build intelligent systems capable of performing complex operations in dynamic, real-world, uncertain, and sometimes hostile, environments.[6] Computer hardware is approaching the estimated speed and memory capacity of the human brain, and is increasing by an order of magnitude every five years. Reference model architectures and software development methodologies have evolved over the past three decades that provide a systematic approach to engineering intelligent systems.[1, 3, 6]

This paper integrates knowledge from all of these disciplines into a computational model of human cortex that is well suited for implementation on next generation computational hardware and software systems.

* james.albus@nist.gov

Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense VII, edited by Edward M. Carapezza,
Proc. of SPIE Vol. 6943, 69430K, (2008) · 0277-786X/08/$18 · doi: 10.1117/12.786262

Our model is predicated on the assumptions that:

1. The brain is first and foremost a control system that is capable of building an internal representation of the external world, and using this representation to make decisions, set goals and priorities, formulate plans, and control behavior with intent to achieve its goals.
2. Perception is a set of processes that gather input from sensors and transform it into a rich, colorful, dynamic, and meaningful internal real-time representation of the external world.
3. What the conscious self perceives is only its internal representation – not the external world. The self interprets its internal representation to be external reality.
4. The internal representation is what the self uses to make decisions, generate plans, and command actions that result in behavior of the body in the world.
5. The self perceives itself at the center of its universe, with the external world represented in an egocentric coordinate frame.

We define a self egosphere as a polar coordinate system with its origin at the centroid of the brain. The self egosphere is divided into right and left hemispheres by a medial plane through the brain. Its north polar axis is coincident with the gravity acceleration vector when the head is held level in an alert position. Our model assumes that the conscious self perceives itself as residing at the origin of the self egosphere facing forward toward zero heading on the horizon.

2. STRUCTURE OF BRAIN

Among the most obvious features of the brain and body is that they are symmetrical about the medial plane. The left side of the brain represents the right hemisphere of egospace and controls the right side of the body. The right side of the brain represents the left hemisphere of egospace and controls the left side of the body. The two sides of the brain communicate through the corpus callosum, the anterior commissure, and many sub cortical structures in the midbrain and spinal cord. The two sides of the cortex form a symmetrical pair of large two-dimensional arrays – one for each hemisphere of the egosphere.

The brain is also partitioned front to back into posterior regions that are primarily concerned with perception, and frontal regions that are primarily concerned with behavior. This partition begins at the bottom in the spinal cord, and continues up the centerline of the cord through the midbrain to the top of the brain along the central sulcus in the cortex.[†] Two-way communications between frontal and posterior cortical regions are accomplished by a multitude of fibers that run fore and aft beneath the cortex. Two-way communications between sensing and acting also take place at many levels within the spinal cord, midbrain, and subcortical nuclei.

It has long been recognized that the entire central nervous system is a hierarchy of computational modules, starting with the sensory-motor centers in the spinal cord, and moving upward through sensory-motor integration centers in the midbrain and basal ganglia, and finally to the sensory-motor levels and echelons in the cortex.[7] More recently, it has become clear that multiple regions in the cortex are organized in hierarchical layers in both the frontal and posterior areas.[2] This is illustrated in Figure 1.

In the frontal cortex, the cortical surface is partitioned into regions that represent tasks and plans at multiple hierarchical echelons, and in multiple coordinate frames. Outputs from higher echelons become inputs to lower echelons. The prefrontal cortex is where high level decisions are made, plans are generated, and priorities are established. Behavioral plans made in prefrontal cortex flow down the behavior generating hierarchy causing computational units in the frontal eye fields to direct the eyes toward regions of interest on the egosphere, and causing units in the premotor and primary motor cortices to move the legs and feet to reach desired locations and maneuver the arms, hands, and finger tips to touch, feel, and manipulate objects of attention. Behavioral commands select procedures to move the body through the environment and search for objects that are needed to accomplish task goals. Status information flows up the behavior generating hierarchy so that higher echelon behavioral centers can monitor the progress of behavior as it evolves, and modify plans to meet unexpected contingencies. At all levels, information flows horizontally in both directions between sensory processing and behavior generating hierarchies.

[†] An exception is the cerebellum, which is a special purpose computer for the kinematic and dynamic transformations that are required for learning, planning, and controlling coordinated muscle activity.

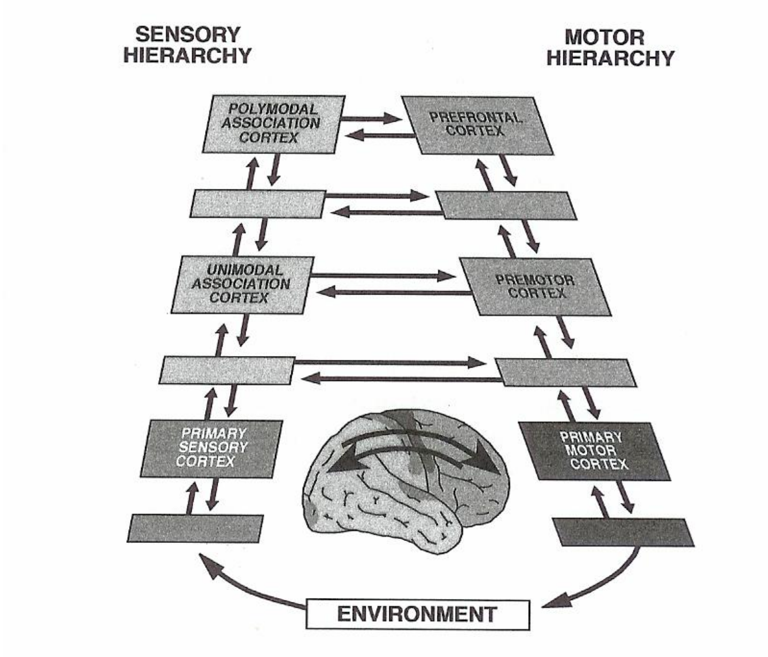

Figure 1. Hierarchical layering of structure and function in both the frontal (motor) and posterior (sensory) regions of the brain. (from [2] by permission)

In the posterior cortex, the cortical surface is partitioned into regions that represent the sensory egospace multiple times, at multiple levels in the sensory processing hierarchy, and in multiple coordinate frames. These regions are interconnected such that, outputs from lower level regions become inputs to the upper level regions. Sensory signals flow up the sensory processing hierarchy to be segmented into patterns, grouped into entities and events, classified, and remembered as situations and episodes. Sensory priorities flow down the sensory processing hierarchy from high-level sensory-motor centers to lower level sensory-motor centers to focus attention on objects that are important to current goals. Sensory priorities cause computational units at all levels in the sensory processing hierarchy to select among alternative processing algorithms for windowing, segmentation, grouping, and classification of sensory signals.

In short, information flows both up and down hierarchies in both frontal and posterior cortex. Information also flows horizontally, both within local neighborhoods and over long distances between posterior-sensory and frontal-motor hierarchies.

There are many subdivisions in the cortex that perform specific functions. The cortex in posterior brain is not only partitioned into a hierarchy of levels, but into separate hierarchies that process input from different sensory modalities. The visual processing hierarchy is located in the occipital cortex. The somatosensory processing hierarchy is located in the anterior parietal cortex. The acoustic processing hierarchy is located in the superior temporal cortex. Outputs from these three unimodal hierarchies are combined in two bimodal association areas: a) visual and somatosensory representations of egospace are combined in the posterior parietal cortex. b) visual entities are merged with acoustic events in the anterior temporal cortex where associations are established between physical objects and written and spoken words. Finally, output from the two bimodal association areas are merged in the region surrounding the junction of the occipital, temporal, and parietal cortices. Output from this region travels to the hippocampus which decides whether short-term memories ought to be consolidated into long-term memory. It is also connected to the amygdala and other parts of the limbic system that assign emotional values such as worth and importance to objects, events, situations, and episodes.

Finally, there are many loops between the cortex and thalamus, as well as between cortical regions within the same level, and between different levels. In the frontal cortex, these loops pass through the striatum and pallidum before passing through the thalamus on the way back to the cortex. In both frontal and posterior cortex, some loops also pass through the cerebellum and various modules in the limbic system.

From this computational architecture emerges the richness of human experience in vision, hearing, feeling, perception, cognition, reasoning, decision-making, planning, and control of behavior that not only reacts to current sensory input, but reasons about the past, anticipates the future, sets long range goals, makes plans, develops tactics, and performs actions designed to achieve the needs and desires of the self.

The overall result of the sensory processing hierarchy is to generate and maintain a rich, dynamic, and colorful internal visual representation of the world that is overlaid with many levels of meaning and emotional value. Top-to-bottom, high level symbolic representations with syntax and semantics are linked all the way down the hierarchy from situations and episodes at the top to individual pixels in sensory images at the bottom. Bottom-to-top, low level pixel representations are segmented into patterns, objects, and events with attributes, state, and relationships that are linked all the way up the hierarchy to trends and classes with attributes that can be inherited and behaviors that can be predicted. Signal from sensor arrays are dynamically linked to objects, agents, relationships, situations, and episodes that have spatial and temporal continuity. Objects are perceived to move about the environment in ways that are subject to laws of physics, and agents can be expected to behave according to rules of social custom.

3. STRUCTURE IN CORTEX

The cortex is a massively parallel structure. The human neocortex is a thin sheet of computational units about 2000 cm^2 (2.2 ft^2) in area and about 3 mm thick. The cortex is remarkably uniform throughout its extent. It consists of six layers of neurons arranged in columns oriented perpendicular to the surface of the cortex. It has long been suspected that these cortical columns perform some kind of computational function in the cortex.[4, 7]

There are two types of cortical columns:

1) microcolumns contain 100 – 250 neurons. These are only 30 μ – 50 μ in diameter (scarcely more than one neuron wide) and about 3000 μ long (i.e., the full thickness of the cortex from layer 1 to layer 6.)

2) hypercolumns (a.k.a. columns) contain 100 + microcolumns in a bundle. These are on the order of 500 μ in diameter and also about 3000 μ long.

A hypercolumn occupies about .2 mm^2 of cortical real estate. Thus, there are about a million hypercolumns in the cortex, and more than 100 million microcolumns. Each of the columns and microcolumns in the cortex are serviced by underlying thalamic nuclei that are connected to the cortex through looping communication pathways. The uniformity of structure in the cortex and the corticothalamic loops suggests that fundamental computational mechanisms may be similar throughout the cortex despite the differences in functional processes that are performed in different regions.

3.1 Communications in the brain

Communications within and between regions in the brain take place via two fundamentally different types of neurons: drivers and modulators.[8] Axons from driver neurons convey specific information that is topologically organized in the form of images or maps. In the visual, auditory, and somatosensory systems, these are maps of egospace. These maps are repeated many times, both within and between levels in the hierarchy.

Axons from modulator neurons select and modulate the computational processes that operate on driver arrays. Modulator neurons have relatively large receptive fields and their axons typically cover large fields of influence with low density and distal synaptic connections on driver neurons. Modulator neurons either do not preserve topological order, or do so very generally.

Our model hypothesizes that cortical modulator axons are essentially address lines that access memory locations where parameters and procedures are stored that operate on driver arrays to focus attention, set segmentation criteria, compute attributes, and classify entities and events. This enables cortical modulator neurons to develop and use dynamic representations to generate predictions that can be compared with driver signals to perform recursive estimation or predictive filtering.

Our model assumes that modulator neurons are able to establish and maintain address pointers that define relationships such as *belongs-to*, *has-part,* and *is-a-member-of* pointers. *has-part* and *belongs-to* pointers link pixels to objects, and vice versa. *is-a-member-of* pointers link entity and event frames to class prototypes. Pointers can also link entities and events in graph structures that describe situations and episodes. In general, pointers set by modulators are able to define graph structures that can represent many forms of symbolic knowledge, including logical predicates, grammars, semantic nets, procedural rules, task skills, scripts, plans, and behaviors.

In summary, our model assumes that drivers communicate specific data (e.g., attributes) that are topologically organized as images or maps, while modulators communicate addresses (e.g., pointers) that access stored memories, control computational processes, and establish relationships. We assume that both attributes and pointers can be stored both temporarily in loops, and permanently in synaptic modifications.

3.2 Computational units in the brain

Our model assumes that the brain stores an internal representation of the external world in arrays of cortical columns. It assumes that correspondence between the internal representation and the external world is established and maintained by hierarchies of computational processes in the cortico-thalamic loops that process signals from sensors. It assumes that these columns and loops are Fundamental Computational Units (FCUs) that, in posterior brain, are organized into sensory-processing and world-modeling hierarchies to transform sensory input into a rich, dynamic representation of the world. It also predicts that, in the frontal brain, FCUs combine with the basal ganglia in the cortico-thalamic loops, to form a behavior generation hierarchy that can transform intentions (i.e., desired goals and priorities) into purposeful behavior.

Our model predicts that a single cortical hypercolumn together with its underlying thalamic nuclei is functionally equivalent to a single FCU.

Our model assumes that each FCU contains three parts:

1) an abstract data structure in the form of an entity, event, or task frame that:
    a) has a location (i.e., a name or address) where it is located in the array of cortical hypercolumns
    b) has membership criteria (that may be variable)
    c) has slots for attributes that define current state, history, and predicted future
    d) has slots for pointers that define spatial, temporal, causal, and semantic relationships

2) a set of computational processes that are able to:
    a) segment spatial patterns of attributes within a spatial receptive field into entity frames
    b) segment temporal patterns of attributes in a temporal receptive field into event frames
    c) set *belongs-to* and *has-part* pointers between entity and event frames at different hierarchical levels
    d) compute entity or event pattern attributes and state and store them in entity or event frames
    e) compare entity or event attributes with class prototypes, and when a match occurs, set *is-a-member-of* pointers to classes
    f) use historical patterns to generate predictions for recursive estimation and planning

3) a set of processors (e.g., synapses, neurons, and circuits in microcolumns and hypercolumns) that implement the above processes.

Our model predicts that the frame data structure for the FCU resides in the cortex, while the computational processes that maintain it are embedded in the circuitry and synaptic properties of the hypercolumn and its underlying thalamic nuclei.

A sketch of the internal structure of our hypothesized FCU is shown in Figure 2. Each FCU contains an entity or event frame with attributes, membership criteria, and pointers. Each FCU also contains computational processes that focus attention, segment lower-level FCUs into those that belong to level(i) FCUs from those that don't, and compute attributes of the level(i) entity or event represented in the level(i) FCU.

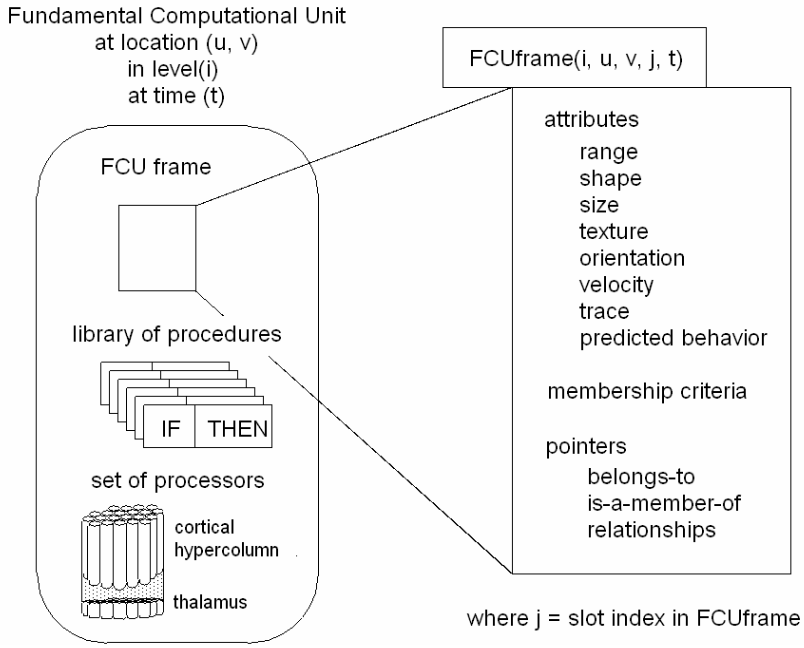

Figure 2. Internal structure of a Fundamental Computational Unit (FCU) consisting of an entity frame, a set of procedures for maintaining the frame, and a set of processors for implementing the procedures. Our models assumes that this FCU is the computational equivalent of a cortical hypercolumn together with its thalamic nuclei.

The data structure for entity frames in FCUs representing hypercolumns in the cortex can be represented as a matrix

FCUframe(i, u, v, j, t)

where
i = level index in the sensory processing hierarchy
u = row index of the FCU receptive field at level(i)
v = column index of the FCU receptive field at level(i)
j = index of slots containing attributes and pointers in the entityframe
t = time

An input/output diagram of a typical FCU in a sensory processing hierarchy is shown in Figure 3.

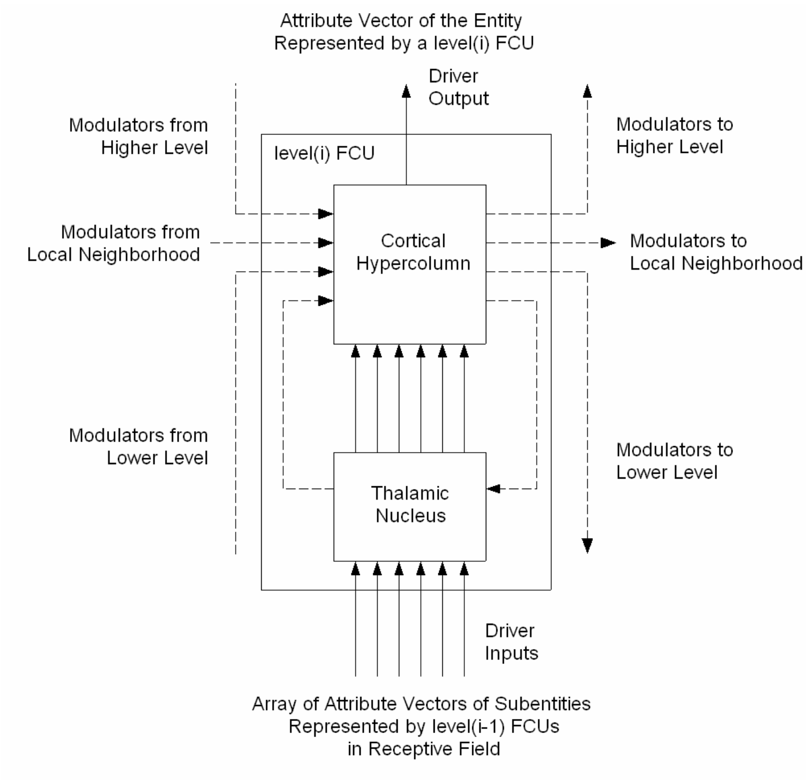

Figure 3. Inputs and outputs of a Fundamental Computational Unit (FCU).

There are eight types of inputs and outputs to and from a typical FCU in the posterior brain:

- *Driver inputs* convey attributes of entity or event frames from lower-level FCUs in the local FCU input receptive field.
- *Driver outputs* convey attributes of the entity or event frame in the local FCU to higher-level FCUs in the output field of influence.
- *Modulator inputs from above* convey commands, priorities, and pointers from above that select processing algorithms, generate masks and windows, provide prioritized list of entities of attention, suggest gestalt grouping hypotheses, and set *belongs-to* pointers in FCU.
- *Modulator outputs to above* convey status of processing results to higher-level FCUs, set alerts regarding regions of the field of view that require attention, and set *has-part* pointers in level above.
- *Modulator inputs from below* convey status of processing results from lower-level FCUs, set alerts regarding regions of the field of view that require attention, and set *has-part* pointers in FCU.
- *Modulator outputs to below* convey commands and priorities to lower-level FCUs that select processing algorithms, generate masks and windows, provide prioritized lists of entities of attention, suggest gestalt grouping hypotheses, and set *belongs-to* pointers in level below.
- *Modulator input from and output to FCUs at the same level in the same map.* These implement lateral inhibition and relaxation processes that fuse information from multiple sources and resolve conflicting evidence.

Modulator signals within a FCU (i.e., within the cortical hypercolumn, and between cortex and thalamus) provide mechanisms for focusing attention, and for segmentation and grouping operations.

4. SEGMENTATION AND CLASSIFICATION

Our model assumes that FCUs at all levels contain frame data structures with slots for attributes and pointers. Within each receptive field at every sensory level, lower level FCUs are segmented into groups that are linked to higher level FCUs. At each level, these frames are linked by *belongs-to* pointers to higher level entity (or event) frames, by *has-part* pointers to lower level entity (or event) frames, and by *is-a-member-of* pointers to class frames. Entity and event FCUs may also be linked together by *relationship* pointers that define situations, episodes, rules, procedures, strings, lists, graphs, and grammars.

Our model assumes that segmentation and grouping processes are embedded in FCUs at each level in the receptive field hierarchy. Thus, for each FCU at level(i), there are segmentation processes that group level(i-1) FCUs into level(i) FCUs based on group attribute criteria. This is illustrated in Figure 4.

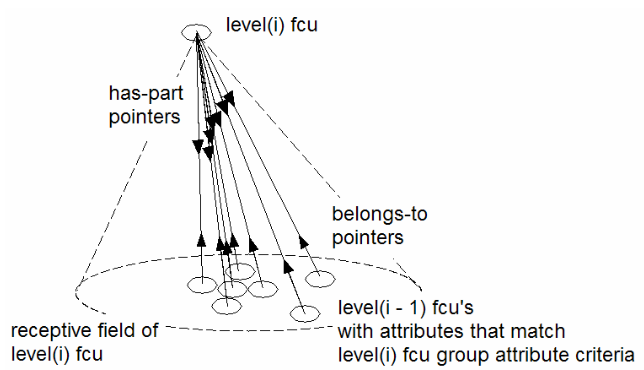

Figure 4. The result of segmentation and grouping. A group of FCUs at level(i-1) are linked by belongs-to pointers to an FCU at level(i). A list of pointers in the level(i) FCU points back to the FCUs at level(i-1).

Group attribute criteria may be based on gestalt properties such as similarity, symmetry, or spatial or temporal proximity or continuity. For example, level(i-1) pixels with similar attributes of color, range, and range gradient, may be grouped into level(i) surface patch entities. Level(i-1) pixels with similar intensity, color, or range gradients may be grouped into level(i) edge-entities.

The result of segmentation and grouping is a set of *belongs-to* and *has-part* pointers that link level(i-1) entities (or events) to level(i) entities (or events.) The set of *has-part* pointers define a membership list for each level(i) entity (or event) that includes all the level(i-1) entities or events that belong to it.

Once level(i-1) entities (or events) have been grouped into level(i) entities (or events), level(i) computational processes compute level(i) entity (or event) attributes. Recursive estimation processes can then compare model-based predictions with sensory observations. Finally, classification processes compare entity (or event) attributes against class prototypes. This establishes *is-a-member-of* pointers that define class membership and enable observed entities and events to inherit class attributes.

A more detailed diagram of the processes and results of segmentation and grouping is shown in Figure 5.

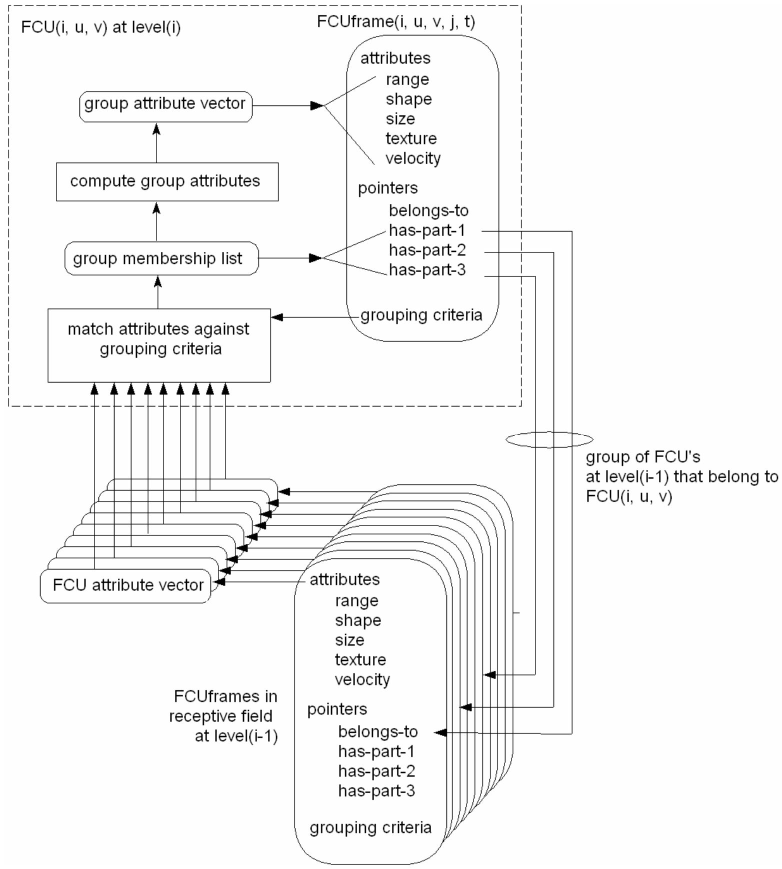

Figure 5. Segmentation and grouping processes compare lower level FCU attributes against upper level FCU grouping criteria. When attributes of lower level FCUs meet higher level grouping criteria, *has-part* pointers are set in the higher level FCU, and *belongs-to* pointers are set in the lower level FCUs.

5. HIERARCHIES OF ENTITIES AND EVENTS

Each level of the sensory processing hierarchy consists of an array of FCUs. The entity (or event) frames in the FCUs are linked by *belongs-to* and *has-part* pointers. Our model asserts that a single entity (or event) frame is embedded in each cortical hypercolumn. This means that *belongs-to* and *has-part* pointers link cortical hypercolumns in entity (or event) hierarchies. An example of a two level entity hierarchy resulting from two levels in the sensory processing hierarchy is shown in Figure 6.

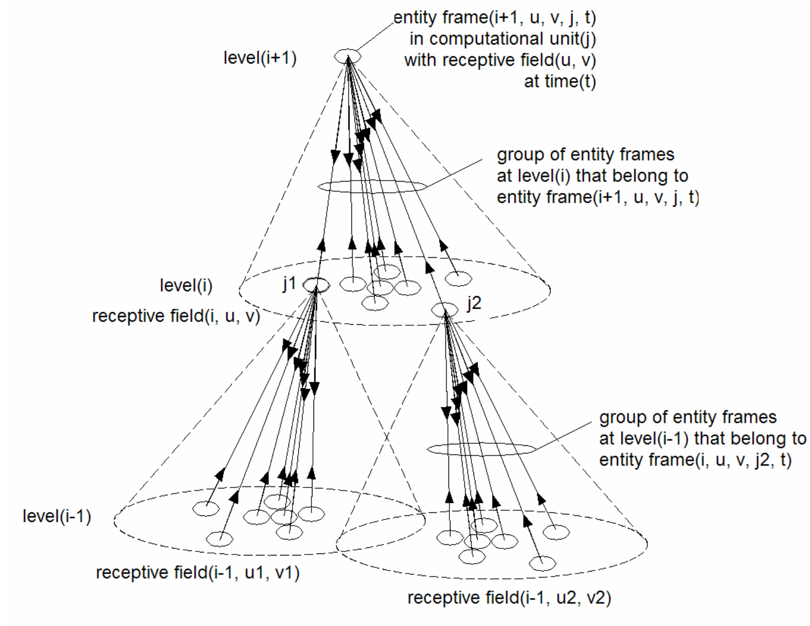

Figure 6. Two levels of segmentation and grouping of entity frames.

The entity hierarchy links pixels to objects and situations that have attributes and can be sorted into classes. Within each Fundamental Computational Unit there are processes that:

1. compute entity attributes such as size, shape, texture, color, intensity, orientation, and motion.
2. generate predicted pixel attributes based on previously estimated entity attributes.
3. compare predicted attributes against observed attributes to perform the function of recursive estimation an adaptive filtering.
4. compare entity attributes with class-prototype-attributes, and set *is-a-member-of* pointers when matches are detected.

The event hierarchy links temporal sequences to events and episodes that have attributes and can be sorted into classes. Event frames are embedded in cortical columns that are linked by *belongs-to* and *has-part* pointers. Event attributes can be compared with class-prototype-attributes and *is-a-member-of* pointers can be established.

In each event-frame, there are processes that:

1. compute event attributes such as start-end points, duration, frequency, time-frequency patterns, and spatial-temporal trajectories.
2. generate predicted signal attributes based on previously estimated event attributes.
3. compare predicted attributes with observed attributes to perform predictive filtering.
4. compare event attributes with class-prototype-attributes, and set *is-a-member-of* pointers whenever a match occurs.

In the auditory processing hierarchy, there are processes that close phase-lock loops, recognize temporal patterns, generate grammars, and assign meaning based on correlation between predicted events based on dynamic representations and observed behaviors of objects and agents in the world.

Thus, for both spatial entities and temporal events, there are FCU processes that maintain pointers that define spatial and temporal relationships in situations and episodes, and provide the basis for grammar, scripts, plans, and patterns of locomotion, manipulation, and language behavior. At all levels, in all of the FCUs, in both sensory processing and behavior generating hierarchies, all of these processes execute in parallel simultaneously.

6. SUMMARY OF THE MODEL

In our model, each hypercolumn and its associated thalamic and other subcortical support nuclei is functionally equivalent to a Fundamental Computational Unit (FCU) that contains a frame data structure, and the computational procedures and mechanisms that generate and maintain this data structure. The FCU frame has slots for attributes and pointers.

In the posterior cortex, our model asserts that these frame data structures are entity or event frames, and that these are linked by pointers in relationships. These relationships can include *belongs-to* and *has-part* links to higher and lower level frames. Relationships can also include class membership relationships, as well as relationships that characterize situation and episodes, and more general relationships such as strings, graphs, rules, and grammars.

In frontal cortex, frames define goals, priorities, tasks, objects, and agents. Pointers define plans, skills, and behaviors. Loops include the striatium, pallidum, and in some cases the cerebellum, where models of body dynamics and kinematics are stored. At the top are high level goals and priorities. At the bottom are individual motor neurons that send commands to control muscles

In the sensory processing hierarchy, frames define entity or event attributes and relationships. At each level in the sensory processing hierarchies, image elements are windowed, segmented, grouped, characterized, and classified. At the bottom, are pixels that represent the receptive fields of sensory neurons. At the top, are abstract data structures that represent complex situations and episodes. These top-level abstract data structures are linked by *has-part* links all the way down to the pixels of immediate experience. Bottom-level pixels are linked all the way up to situations and episodes by *belongs-to* links.

Our model reasons about the brain from the top down. Our goal is to understand both form and function. Our model suggests an anatomical and computational architecture by which sensory signals can be transformed into an internal representation of the external world that is both visually rich and overlaid with meaning and purpose. It suggests how the neural structures in the brain can perceive the world as a rich, dynamic, colorful place filled with people, places, and events that can be assigned worth, and linked to emotional values.

Our model assumes that the sensory processing cortex contains two types of hierarchy:

1. a receptive field hierarchy defined by the connectivity of the driver neuron pathways, and
2. an entity or event hierarchy defined by segmentation and grouping algorithms that detect patterns of activity within receptive fields and cluster subentities into entities (or subevents into events.) Segmentation and grouping criteria are gestalt properties such as proximity in space or time, similarity of attributes, symmetry, and smoothness.

These two hierarchies are embedded in the same neural architecture, but they are distinct in form and function.

Our model predicts that:

1) abstract data structures are stored cortical hypercolumns in the form of patterns of neuronal spiking in cortical microcolumns. These are maintained during immediate experience by local recursive filter loops. They are maintained in short-term memory by the neural equivalent of finite-state automata, and in long-term memory by physical modifications of synapses. Our model predicts that each hypercolumn contains a single entity frame, and that microcolumns within the hypercolumn are tuned to compute attributes and state variables, and to establish and maintain pointers.

2) modulator neurons compute addresses that can elicit patterns of activity that generate masks and windows and select segmentation and grouping hypotheses. The exchange of modulator signals between hypercolumns at different levels in the sensory processing hierarchy establish *belongs-to* and *has-part* pointers between entity and event frames at various levels.

3) modulator pathways between hypercolumns at different hierarchical levels travel horizontally just under the cortex, and do not flow through the thalamus. These are capable of establishing reciprocal pointers that define

many types of relationships, such as class membership, situations, episodes, rules, programs, and behaviors. They can also enable belief systems, logical reasoning, grammar, semantics, and language.

4) *belongs-to* and *has-part* grouping relationships enable images to be interpreted in terms of surfaces, objects and groups that are recognizable and linked to classes from which they can inherit prior expectations and emotional values. Thus, we perceive the world not as an array of pixels, but as objects, events, situations, and episodes that are subject to the laws of nature; with spatial and temporal continuity, purpose, and meaning. This network of pointers solves the "symbol grounding problem" and the "binding problem."

5) The massively parallel nature of this architecture enables the network of pointers to be established quickly and updated rapidly. Our model predicts that pointers are established or updated between any two hierarchical levels in the brain within about 10 ms, and all the way from bottom to top of the unimodal hierarchies in about 100 ms. With each saccade of the eyes or motion of the fingers, low-level pixels and list entities are re-linked to high-level objects and situations that have spatial and temporal continuity.

In the future, we hope to extend our model to address how the brain uses its internal world representation to reason about the past, to contemplate the future, to make decisions and plans, to set goals and priorities, and to respond to conditions in the environment in a manner designed to achieve its goals.

7. REVERSE ENGINEERING THE BRAIN

There are about a million hypercolumns in the human cortex. Thus, it would take about a million FCUs to build a full-scale computational model of the human brain. Because each of these computational units has a similar internal structure, this model is well suited to the massively parallel architecture of modern supercomputers. A modern supercomputer has hundreds of thousands of computational cores that provide computational capacity on the order of 300 teraflop (3×10^{14} floating operations per second). If we divide that up between a million FCUs, we get 3×10^{8} flops per FCU. For real-time operation at 200 compute cycles per second, we obtain about 1.5×10^{6} flops per FCU per compute cycle. This seems adequate to perform the operations we envision taking place in our FCU. For some types of operation, special purpose hardware may enable significant increase in computational capability.

Of course, modeling computational units is only part of the problem. The FCUs need to be connected such that outputs from FCUs in one compute cycle become inputs to other FCUs in the next compute cycle. While is seems likely that a modern supercomputer has sufficient computational power to emulate the functionality of the FCUs, communication between FCUs may be more problematic. There are many unanswered questions. For example, how many FCUs must communicate with each other? How often? How much information needs to be communicated?

Fortunately, not everything in the brain connects to everything else. All axon connections are point to point. Communications between neurons occurs over what is essentially a publish-subscribe network. Each neuron publishes information on an axon that carries it to a finite set of subscribing synapses on target neurons.

One approach might be to adopt a messaging system such as the Neutral Messaging Language (NML)[3] developed at NIST for communications between modules in real-time intelligent control systems. NML is a publish-subscribe communication architecture that could be extended to become a *Neural* Messaging Language. The connectivity specified in the NML configuration file would need to be informed by neuro-anatomical data. Receptive fields and fields of influence of both drivers and modulators would need to be reflected in the publish/subscribe configuration files. Driver inputs and outputs are arrays of attribute vectors. Modulator inputs and outputs are addresses that can be used to access data, page memory, set priorities, select procedures, and set pointers within and between FCUs. At the end of each computational cycle, each FCU must write to its NML output buffers. At the beginning of the next cycle, each FCU must read from its NML input buffers. In the period between compute cycles, NML must move messages from publishers' output buffers to subscribers' input buffers. To mimic real-time performance in the brain, the combined computation-communication cycle should repeat about every 5 ms.

At present, we have no estimate of how big a NML configuration file must be, or how long it might take for NML to move all the messages from input to output. This will depend on the hardware implementation. It is an engineering question that will have an empirical answer. Of course, there are other message passing systems that have been developed in the robotics and computer science communities. Some of these may be more appropriate than NML for modeling axon connectivity.

The bottom line is that there is growing reason to believe that reverse engineering the human brain may be a feasible scientific and engineering goal. Almost certainly, any roadmap designed to achieve that goal will focus first on reverse engineering the structures in the brain responsible for visual perception and image understanding. Much is known about the structure and functioning of the first few levels in the image processing hierarchy. Understanding the role of the cortical column and its interactions with its underlying thalamic nuclei in the visual cortex may provide a Rosetta stone for understanding the rest of the brain.

ACKNOWLEDGEMENTS

This paper was written as part of the official duties of the author as an employee of the U. S. government. It is therefore not subject to copyright.

REFERENCES

[1] Albus, J. S., Meystel, A., [Engineering of Mind: An Introduction to the Science of Intelligent Systems] Wiley, New York, (2001).

[2] Fuster, J. M., [Cortex and Mind: Unifying Cognition], Oxford University Press, Oxford, U.K., (2003).

[3] Gazi, V., Moore, M.L., Passino, K.M., Shackleford, W.P., Proctor, F.M., and Albus, J.S., [The RCS Handbook: Tools for Real Time Control Systems Software Development], John Wiley & Sons, NY, (2001).

[4] Hubel, D.H., and Wiesel, T.N., "Receptive fields and functional architecture of monkey striate cortex." 195: 215-243 (1968)

[5] Kandel, E., Schwartz, J., Jessell, T., [Essentials of Neural Science and Behavior], McGraw-Hill, New York, (1995).

[6] Madhavan, R., Messina, E., Albus, J. S., [Intelligent Vehicle Systems: A 4D/RCS Approach], Nova Science Publishers, New York, (2007).

[7] Montcastle, V. B., "The columnar organization of the neocortex," *Brain* 120 (Pt. 4) 701-722 (1997).

[8] Sherman, S. M., and Guillery, R.W., [Exploring the Thalamus and It's Role in Cortical Function], MIT Press, Cambridge, MA, (2001).