

Chapter 16

Overview and Reflections

The goal is to understand in as much detail as possible the biophysical mechanisms and systems of the human brain that account for our cognitive capacities. I presented a partial list of basic cognitive tasks that confront us in our day-to-day activities and that we are able to perform with apparent natural ease. A credible theoretical model of the cognitive brain must specifically address these ecologically relevant tasks and provide evidence of its operational competence in the appropriate cognitive domains. At the same time, the fundamental biophysical and structural assumptions of the model must be plausible in the light of current neurophysiological and neuroanatomical knowledge.

Although the brain model I have proposed does not cover the full range of cognitive processes in each modality, it does explicate a number of plausible neuronal mechanisms serving a variety of essential functions. These mechanisms and systems were shown to be competent for tasks that are critically important in human cognition. In addition, the same putative brain mechanisms predict classical visual illusions and other anomalous perceptions as epiphenomena of their normal operating principles. Moreover, a large and diverse body of experimental and clinical findings can be explained by the biophysical and structural properties of these mechanisms.

The principal functions mediated by the hypothesized mechanisms can be subsumed under the following broad headings: learning and long-term memory, short-term phasic representation, visual cognition (including imagery), semantic processing, planning and composing behavior, and motivation. Neurons are the primitive biological units of which the entire model system is composed. A block-flow diagram of the integrated cognitive brain system is shown in figure 16.1.

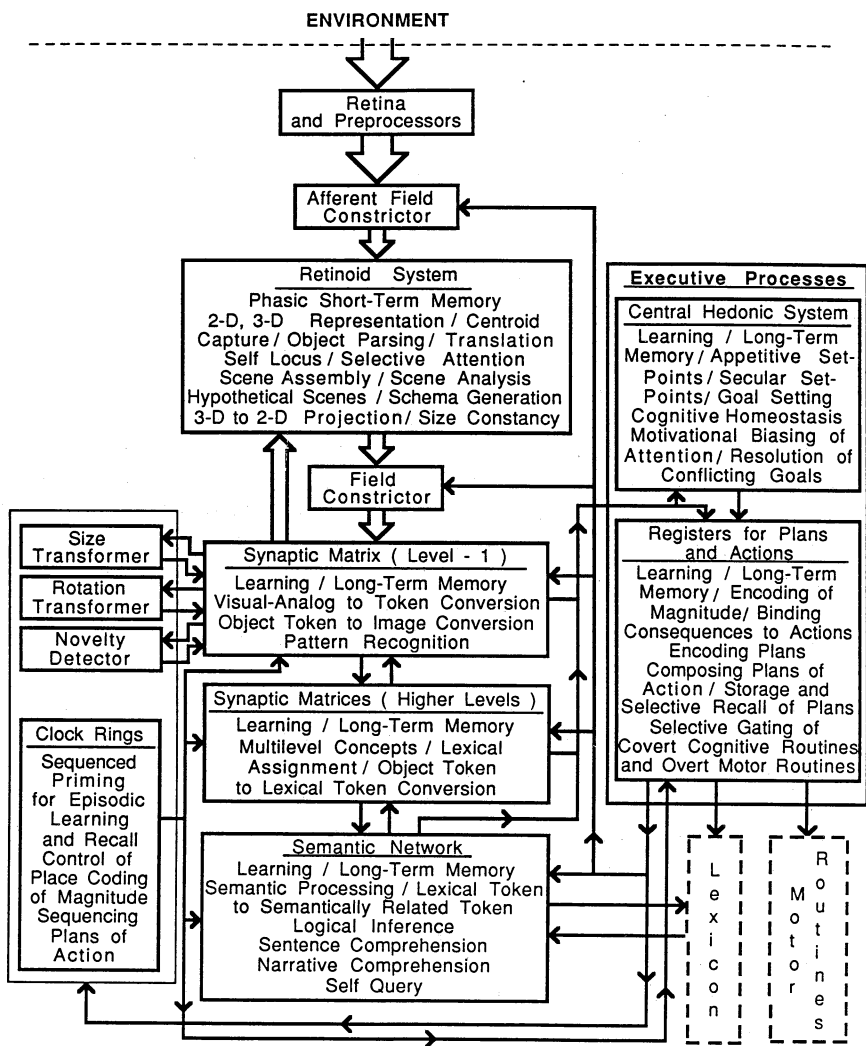


Figure 16.1
Block-flow diagram of the cognitive system.

Learning and Long-Term Memory

The biophysical aspect of learning is a long-term modification (increase) of synaptic transfer weights (ϕ) in selected synapses of adaptive neurons. Learning will occur if and only if coactivation of a pre- and postsynaptic cell results in postsynaptic activity sufficient to result in a reaction between an axon transfer factor (ATF) and a free dendrite transfer factor (DTF) within the postsynaptic dendritic matter. During an ATF-DTF reaction in an adaptive cell, ϕ increases occur only at active synapses, and the peak magnitude of ϕ at each synapse is limited by the amount of DTF that can be utilized at the local receptor region. Whenever learning occurs, ATF makes a small, fixed contribution to the transfer weight of its locally active synapse, whereas a limited store of DTF is distributed over all active synapses and makes a contribution to each that is, up to the local saturation limit, inversely proportional to the number of concurrently active axonal inputs. This latter property causes intrinsic quasi-normalization of the synaptic weight distribution on an adaptive cell and is critical for effective pattern learning and subsequent recognition. These assumptions are reflected in the basic learning formula represented in equation 2.3.

The distribution of synaptic transfer weights on each adaptive cell constitutes its latent memory for the stimulus that it learned (its pattern of axonal inputs). We can think of such a cell as having been tuned to a particular stimulus, which will be more effective than any other in activating the adaptive cell. For this reason, I call a neuron that has been synaptically modified in the course of learning a labeled line.

Another point should be emphasized: learning will not normally occur simply on the basis of coactivation of presynaptic cells and adaptive postsynaptic cells; the concurrent postsynaptic activation must be high enough to exceed a threshold for effective ATF-DTF interaction. I have hypothesized that this is accomplished by auxiliary excitatory priming of adaptive neurons during heightened arousal (Trehub 1975a, 1977) and, for episodic learning, by additional priming from the autaptic cells of a clock ring (Trehub 1983).

Short-Term Phasic Representation

The ability to capture arrayed patterns of excitation or single neuronal events in short-term memory is necessary for a variety of cognitive processes. This capability is assumed to be mediated by autaptic neurons. Cells of this kind have one or more of their axon collaterals in

feedback synapse with their own dendrites or cell body. If there is sufficient sustained excitatory bias on such cells from an auxiliary source, a transient stimulus will cause autaptic cells to continue firing even after input from the initiating stimulus has stopped. If the auxiliary bias is removed, however, excitation from its own recurrent axon collateral is insufficient to sustain spike discharge, and the autaptic cell stops firing. Autaptic neurons are key elements in a number of cognitive modules, such as the retinoid system, clock circuits, latching circuits, and mechanisms for mapping stimulus magnitude to labeled lines.

Visual Cognition

The cognitive brain has a remarkable capacity to learn the ecologically significant properties of the visual world and to model these properties within its neuronal structures. Humans are not only able to learn, represent, and analyze environments that they have directly experienced but also are able to construct and analyze internal models of imaginary and novel environments. In my theoretical model of the cognitive brain, this ability is instantiated by the normal operation of a subsystem of explicated neuronal mechanisms, chief among which are the synaptic matrices and the retinoids.

Synaptic Matrix

The synaptic matrix is the basic neuronal mechanism for learning, memory, and imagery. Its adaptive elements are called filter cells and mosaic cells; synaptic modification in these neurons occurs according to the constraints of the learning formula. The synaptic matrix at the first level of visual learning receives its input in the form of a parsed excitation pattern from the retinoid system and maps the pattern as an exemplar of its kind onto a discrete output cell called a class cell, which constitutes a physical symbol or token of the exemplar it has learned. Recurrent axon collaterals of class cells synapse on dendrites of the mosaic cell array in the synaptic matrix. When a class cell is discharged, it evokes an analog image of the exemplar with which it was associated during the course of learning. Thus the brain is able to recall a visual image of a learned object or scene simply by the firing of its class cell token.

Input patterns to the synaptic matrix can be transformed in size and angular orientation by the action of two integrated neuronal circuits, a size transformer and rotation transformer. This enables the system to learn a single exemplar in a variety of sizes and orientations

or to transform a given stimulus to test its match against patterns that have been learned.

Higher-level synaptic matrices organize hierarchical conceptual representations by adaptively forming selective synaptic links among neuronal tokens that signal levels of increasing abstraction. For example, class cell tokens of a particular cat and a particular dog may be mapped to a single token that signifies *pet*; tokens of a car and a television set may be mapped to a token that signifies *artifact*; and tokens of both *pet* and *artifact* may be mapped to a yet higher-level token signifying *property*. The normal operating characteristics of the model provide automatic backward chaining among hierarchically organized neuronal tokens down to the first level of sensory image evocation on the mosaic cell array. This means that when a class cell token of a high-level concept is discharged, a neuronal representation (image) of its semantic content is also evoked.

In addition to the ability to learn and recall static semantic representations, the synaptic matrix is capable of learning sequences of tokened experience that are related and rendered meaningful by their temporal contiguity within particular time frames (episodic representations). Episodic learning and recall in which temporal excursions of remembrance can be brought under motivational guidance is achieved by sequential priming of filter cells and class cells. Clock rings and recall rings are the auxiliary mechanisms that provide the priming sequences for episodic representation.

Retinoid System

A retinoid is a reticulated sheet of autaptic cells each joined to its neighbors by excitatory and inhibitory interneurons. It receives retinotopic input from visual afferents and from images on the mosaic cell array of the synaptic matrix. It serves as a visual scratch pad with spatially organized information (patterns of autaptic cell activity) stored as short-term memory. The retinoid also has an important dynamic property in that selective excitatory biasing of its interneurons by shift control cells can move a "captured" pattern of autaptic discharge to any position on a retinoid surface. The complete retinoid system consists of many interconnected retinoid layers with specialized functions. Within subsystems, patterns of autaptic activity on one layer can be projected to other layers and combined with previously captured patterns to compose extended and complex retinoid displays. Output from the central region of the retinoid system is projected to the mosaic cell array of the level-1 synaptic matrix for learning, recognition, or analysis.

Retinoids are spatially organized around a central point of reference, the normal foveal axis. This is the key axis of egocentric reference and corresponds to the line of sight when the eyes and head point straight ahead and the shoulders are square with the upright body. I hypothesize the existence of a specialized retinoid layer, which I call a self-locus retinoid. Within this layer, a spatially compact region of sustained autaptic cell discharge is centered on the normal foveal axis and defines the origin or home position of the self-locus. Selective activation of shift control cells can cause heuristic excursions of self-locus excitation from its home position to other regions of retinoid space. Since spatially focused excitation from the self-locus retinoid can be projected to corresponding coordinates on other retinoid layers, local excitatory priming of the corresponding regions will be induced. This process provides one of the neuronal means for selective attention.

Excursions of the self-locus also provide another way for mediating selective spatial attention. Shifting the heuristic self-locus to a particular region of retinoid space entails a set of shift commands on horizontal and vertical shift control cells. If the self-locus has already settled in a retinoid region, then any subsequent visual target in that area can, by a simple reversal of the shift commands already specified for the self-locus excursion, be immediately translated to the normal foveal axis, where it will be projected to the synaptic matrix for recognition.

The retinoid surface is quadrantly organized so that autaptic cell activity on right and left hemifields and top and bottom hemifields can be independently summed. This property, together with the auxiliary mechanisms of the afferent field constrictor and the hemifield disparity detector, enable the system to parse complex stimulus patterns and to position the centroids of their parsed retinoid representations on the normal foveal axis. As a result of this intrinsic normalizing capability, the number of exemplars of an object that must be learned and stored in the memory of the synaptic matrix for effective recognition is greatly reduced.

The complete retinoid system can be roughly divided into two major subsystems: a 2-D subsystem consisting of a separate monocular module for each of the eyes and the 3-D retinoid, which receives input from each of the 2-D retinoid modules. It is characterized by having multiple retinoid layers corresponding to the egocentric distances of visual targets in depth. The retinoid layers in this module are arrayed along the Z-axis and are called Z-planes. Thus the monocular subsystem represents objects on X,Y retinoid coordinates, and the binocular (3-D) subsystem represents objects on X,Y, and Z coordinates.

dinates. In addition, heuristic excursions of the self-locus can be made along the three axes of the 3-D retinoid; its coordinates can be sensed, learned, and stored in a synaptic matrix where they can be adaptively associated with any object that it has targeted. This capability allows us to recall where an object was located as well as to know what the object was.

An important structural feature of the 3-D retinoid is the arrangement of its axonal projections to the mosaic cell array of the synaptic matrix. There are spatiotopic projections from each Z-plane to the mosaic array. But from the nearest Z-plane to the farthest, there is a corresponding progressive expansion of axonal projections to the target mosaic cells. The result is an increasing magnification of an object's represented size on the mosaic cell array as the distance of an object increases. This connectivity property can account for the phenomenon of size constancy.

Among the many visual-cognitive functions served by the retinoid system are the following:

1. Parsing objects in complex visual environments.
2. Constructing 3-D visual representations of objects and scenes.
3. Locating a representation of the self within a represented environment.
4. Representing the paths of moving objects and paths of self-excursion in complex environments.
5. Effecting selective shifts of focal attention.
6. Performing geometric and relational analysis of veridical and/or hypothetical objects and scenes.
7. Modeling situational schemas.

Although the retinoids are short-term memory mechanisms, information that is represented in their excitatory patterns can be selectively captured in a synaptic matrix and stored in long-term memory.

Semantic Processing

A simple neuronal elaboration of the basic synaptic matrix enables it to operate as a network for semantic processing. In this mechanism, a set of autaptic cells captures lexical tokens in short-term memory. When a sentential proposition is learned, the activity of these word cells is sustained during the temporal span of each sentence and provides input to the mosaic cells of the matrix so that subject and predicate are adaptively linked through synaptic modification. After learning has occurred, the semantic network can be queried by the discharge of a mosaic cell or a class cell. The discharge of any particu-

lar inosaic cell selects a token of a word or phrase taken as the subject of a sentence and evokes its associated predicate tokens. The discharge of any particular class cell selects a token of a word or phrase taken as the predicate of a sentence and evokes its associated subject tokens. A query of the first kind defines a selected subject in terms of its learned predicates; a query of the second kind infers a subject on the basis of selected predicates.

The putative neuronal mechanism for semantic processing is not only able to respond appropriately (logically, sensibly) to queries related to single sentences, it can also make sense out of combinations of logically related sentences. Moreover, it can deduce chains of logical implications that are only implicit over a body of lexically based knowledge. The semantic network is capable of narrative comprehension, including appropriate pronominal reference, and can deal effectively with negation.

Self-query can be generated by endogenously evoked discharge of lexical tokens within the semantic network. In this way, the neuronal mechanism can monitor the semantic and logical implications of lexical communications in an on-line fashion without mediation or prompting from an outside source. Communications subject to semantic processing can be self-generated sentences (inner speech) as well as the lexical productions of others. By means of backward chaining of class cell tokens, lexical events can be bound to neuronal representations (images) of events in the real world. Inferences drawn by the semantic network can provide information on which to base reasonable plans of action. These plans, in turn, can be expressed and learned as new tokened representations (lexical productions) of action schemes that are subject to elaboration and interpretation within the semantic network.

Planning and Composing Behavior

A subsystem within the executive processes of the cognitive brain consists of a semantic network, neuronal mechanisms for encoding physical magnitudes, synaptic matrices for binding consequences to actions, and adaptive neuronal registers for learning, storing, and selectively executing action sequences. These interrelated mechanisms provide the biological substrate for composing plans and controlling the particular sequences of behavior called for by each plan. Their neuronal properties and component circuits entail the same basic operating principles employed in the other modules of the cognitive brain.

Self-query is an essential operation in the process of composing a plan. Having the ability to recognize particular situations and to recall the specific behavioral affordances bound to each situation, the system is able to learn elementary tokened propositions of the form *A gets B* where *A* represents a situation and *B* represents a different situation or goal that can be achieved by the affordances inherent in *A*. Starting with a desired goal (say *G*), the self-query "What gets *G*?" evokes the necessary situation-action (say *H*) that obtains *G*. Iteration of this operation back to a desired starting situation completes a chain of neuronally tokened situation-contingent actions that constitute a plan. When plans are composed, they are stored as long-term resources in the synaptic weight distributions of neuronal registers, where they can serve the efficient pursuit of goals and the solution of diverse ecological problems. Given appropriate motivation, a plan is selectively initiated when a current situation is recognized as similar to an earlier situation in which the particular plan achieved the desired goal (either in fact or in imagination).

Motivation

Motivation, the core function among the executive processes, is biologically instantiated and developed by a complex of specialized interoceptive sensors and neuronal circuits that I have called the central hedonic system. Within this system, I assume two principal kinds of homeostatic subsystems, designated as HS-I and HS-II. In HS-I, the loop of events serving to hold a physiological state near an optimal set point does not require the activation of a neuronal token of an external referent. An example of HS-I is modulation of cardiac output to hold blood pressure within a normal range. HS-II, which I call cognitive homeostasis, is exhibited by all of the processes that depend on the discharge of one or more referential tokens to complete a homeostatic loop that can correct a deviation from an internal set point. Processes of this type may be entirely covert or may require instrumental behavior to close the homeostatic loop. An example is obtaining a glass of water to satisfy thirst.

The affective brain states of pleasure and displeasure are intimately related to individual motivation and are assumed to be modulated by homeostatic events. Experimental findings indicate that a necessary condition for the experience of pleasure (activation of its cellular substrate) is a deviation from a homeostatic set point in the brain. This, together with some action or event that tends to restore equilibrium to the perturbed system, seems sufficient to evoke pleasure. When homeostasis is achieved, the experience of pleasure dissipates.

Secular set points are distinguished from those that directly serve vegetative and appetitive homeostasis. A secular set point is an enduring brain state that represents an imagined and desired worldly situation—a personal goal. As long as the imagined neuronal representation of a goal is not effectively matched by the perception of its corresponding events in the real world, the homeostatic system of which it is a part remains in disequilibrium. This is the neuronal mechanism that initiates goal-directed behavior. Plans and actions that lead to a situation that matches the goal are selectively activated and are said to be motivated by the goal.

A specialized set of autaptic cells within the central hedonic system contains the discrete neuronal tokens of secular goals. In order for one of these cells to be instantiated as an active goal in the networks of the brain, it must be stimulated by the token of an imagined goal situation and receive a priming margin of excitatory input from the sum of associated adiant events sufficient to overbalance the inhibitory input from the sum of associated abiant events. Under these conditions, an autaptic token of a secular goal will fire and sustain its activity as long as the margin of input from adiant sources over input from abiant sources is great enough. The continued discharge of a goal token represents the departure from the set point (disequilibrium) of a specific secular goal and signals the need to construct an appropriate plan of action for achieving the goal. Recognition of a realized goal sends an inhibitory input to the goal token, squelching its autaptic activity and restoring equilibrium.

An enormous number of objects and events stimulate our senses, but only a relatively few elicit significant responses. We exhibit a selective sensitivity to those that must be discriminated in order to reach a goal and ignore the rest. This kind of selective perceptual sensitivity is automatically induced through associative excitatory priming of discrete filter cells (in detection matrices) by the activated tokens of autaptic set points (goals) that happen to be in a current state of disequilibrium.

The motivational system must resolve the problem of multiple current goals. One of the putative brain mechanisms in the system resolves competition among many different energetic goals on the joint basis of the current relative strength (demand) of each goal and the presence or absence of affordances needed to continue on a path to each goal. In this model, control of behavior is captured by the strongest tokened motive having the supporting bias of perceived available affordances.

Neuronal Capacity

If we accept the hypothesized neuronal system as a strong candidate model of the cognitive brain on the grounds of neurophysiological plausibility and functional competency, can it pass the test of accommodation within the capacity limitations of the brain? As it learns more about the world during normal ontological development, will performance be prematurely impaired because the necessary long-term store of world knowledge too quickly exhausts the available neuronal resources of the brain? No definitive answer can be given, but some reasonable estimates suggest that the putative mechanisms are viable over a human lifetime.

Let us assume an average life span of 80 years. As a rough estimate, let us say that because of the ubiquitous novelty of new-found experience and a generally high level of arousal, the young child learns several new things from external and internal sources every minute of a 16-hour waking day. (Of course, the child's waking day is usually less than 16 hours, but for the sake of simplicity, let us posit 16 hours per day for every day of one's life as the period in which learning occurs.) Assume that there is a general decline in the number of things learned as one ages so that, on average, during an 80-year span, something is learned every 2 minutes (0.5 events per minute) when one is awake. I think it would be fair to characterize this as very active learning—perhaps unnaturally active, but in the context of the analysis it can only strengthen the conclusion. We can now calculate $80 \text{ (years)} \times 365 \text{ (days/year)} \times 16 \text{ (hours/day)} \times 60 \text{ (minutes/hour)} \times 0.5 \text{ (things learned/minute)} = 14,016,000$ as the estimated number of things learned over an 80-year lifetime. (Landauer 1986 provides a different approach to the estimation of lifetime memory.) If each learned event were represented by a discrete filter cell—class cell (token) couplet in lifelong memory, then 28,032,000 neurons would be needed to store the accumulated information.

I have proposed an essentially deterministic cognitive model, but there is considerable evidence that underlying random events and stochastic processes can influence some aspects of brain activity (Gersuni 1971, Trehub 1971a, 1971b, 1973, 1975b). Suppose that when a thing is learned, random perturbations affecting the thresholds of unmodified filter cells often result in the redundant storage of a single learned event. If this occurs, then the number of available (unmodified) adaptive cells in the brain will be "used up" more rapidly. On this account and in the interest of making a more conservative estimate, let us increase by a factor of 10 the number of neurons needed to store accumulated information (memory). Now the cognitive brain

will require approximately 280,320,000 cells to store the information acquired over a lifetime. But even if we accept this increase, a much larger part of the brain's neuronal resources will still be available for other functions. On the basis of cell counts, the estimated number of neurons in just the cortical mantle of the human brain ranges from 10 billion to somewhat more than 16 billion. If we assume that, on average, there are approximately 13 billion cortical neurons, the number dedicated to accumulated memory over a full lifetime represents only 2 percent of the total number of neurons in the cortex of the human brain. This means that even under a conservative estimate that allows for random redundancy of memory storage but does not allow for cellular forgetting and reuse, over 12.5 billion cortical neurons are still available for other processes after all memory requirements are met. Chief among the other neuronal structures are the sensory afferents, retinoids and mosaic arrays, accessory mechanisms, simple feed-forward and feedback lines, the central hedonic system, and adaptive interfaces between major processing modules in the various sensory modalities as well as the action-command system. I believe that these structures can be effectively accommodated without straining the resources of more than 12 billion cortical neurons.

Synaptic Capacity

Cytological evidence suggests that there are more than enough synaptic connections in the neocortex to satisfy the structural requirements of the proposed model of the cognitive brain. It is estimated that the number of synapses on individual cortical neurons ranges from fewer than a few hundred to over 100,000 for the larger pyramidal cells. An estimate of 2000 is often given as the average number of synapses per cell. On this basis, we can say that there are about 2.6 trillion synaptic junctions in the human cortex.

In the case of adaptive cells that accept relatively few axonal inputs, there is clearly no problem about the number of synaptic sites that must be accommodated on the dendrites of each cell. However, hypothesized mechanisms that are characterized by very large numbers of parallel axonal inputs on single neurons require a closer look. For example, each filter cell in the level-1 synaptic matrix receives from the mosaic cell array a large number of visual afferents from retinoid representations of retinal images. Retinotopic excitation is carried in parallel to the higher centers of the brain through a bundle of ~1 million optic nerve fibers (Kuffler, Nicholls, and Martin 1984).

If effective performance of the proposed model were to require a correspondingly large array of axo-dendritic synapses on each filter cell in the synaptic matrix, it would not be a physiologically credible model; however, only a small region around the normal foveal axis (10 degrees or less of visual angle) needs to be projected to the synaptic matrix for effective object recognition. Even this limited number of visual projections can be markedly reduced without causing significant stimulus degradation if the afferent neurons sample retinal activity at an increasing interval as a function of retinal eccentricity. In addition, should the need for a great abundance of synaptic sites require a dendritic surface larger than that provided by a single cell, two or more cross-coupled filter cells converging on a common class cell would serve the purpose.

Combinatorial Explosion?

The fact that a distinct rigid object can undergo a vast number of fortuitous transformations in 3-D space poses one of the seemingly insurmountable challenges to an understanding of human visual pattern recognition. Each object can project to any one of a great many retinal locations; at each location, it may be represented in any one of a great many different sizes; at each size, it may be represented at any one of a great many angular orientations on the 2-D frontal plane; at each angular orientation on the frontal plane, it may be represented at any one of a great many angular orientations on the 3-D depth plane. Any effort to capture (learn) each representation of an object over all its possible transformations is thwarted by what has been dubbed a combinatorial explosion—a demand for storage beyond any conceivable brain capacity. If the brain cannot store all the diverse representations of discrete objects, how can it recognize them given their transformational vicissitudes in the natural world?

The visual-cognitive system that I have hypothesized solves the problem of combinatorial explosion in three principal ways. First, the range of proximal neuronal representations of any discrete object that the detection matrices might have to discriminate is greatly reduced by the action of specialized mechanisms and structures (size transformer, size constancy projections from Z-planes, rotation transformer). These can perform transformations of the size and angular orientation of any given visual representation so that its shape can be tested for a match against the shapes of previously learned patterns having different sizes and orientations. Second, before an object is tested for recognition or is learned, its retinoid representation for

projection to the mosaic cell array of the synaptic matrix is automatically translated to a standard egocentric location (centroid on the normal foveal axis). This means that the brain does not have to store multiple representations of an object to cover all of its possible retinal locations. Third, objects are normally learned only when they are detected as novel. This too reduces the memory load on the brain.

For some kinds of objects, particularly those that imply immediate danger, it is critically important to recognize the object and make an appropriate response as quickly as possible. It is likely that objects of this kind will be learned at a number of different orientations to eliminate the need for time-consuming spatial transformations before they can be recognized. For example, in chapter 10 it was demonstrated that when different objects (though of roughly similar shape) were each learned at just two angular orientations separated by 22 degrees, presentations of the same objects that differed by 11 degrees in orientation from either stored representation were always recognized correctly. Thus, it seems that learning a few exemplars over a range of possible transformations can insure reasonably accurate as well as rapid response.

The Pragmatics of Cognition

The cognitive brain is a pragmatic and opportunistic organ, selectively favored in its evolutionary development to the extent that it has been able to contribute to the solution of ecologically relevant problems. Excluding the physical limitations of our sense organs, there are few a priori constraints on what will constitute the contents of our cognitive apparatus. Operating characteristics of the brain model I propose dispute the commonly held notion that we "carve the world at its joints" before objects in the visual environment are learned and stored in memory. Instead, we first learn whatever novel parts of the extended world happen to be captured within the visual-afferent aperture by the centroid-based parsing mechanism at a time when arousal is sufficiently high. Since arousal is typically increased by energetic needs and motives (homeostatic imbalances in the central hedonic system), which govern our actions, the parts of the world selected for learning will naturally be those that are experienced together with high motivation and goal-directed action. In short, among innumerable possible partitions, we tend to learn roughly those pieces of the visual world that have ecological utility.

I do not imply that we do not carve the world at its joints but rather that this is a later perceptual-cognitive process. Primitive learned parsings can be imaged and projected to the retinoid system for as-

sembly and analysis. Complex retinoid patterns can be decomposed into components that are projected back to the synaptic matrix, where they can be learned and mapped to class cell tokens. If such abstracted objects were to be imaged, we would take them to be the "real" parts of the world. But the boundaries of decomposition would be drawn, I believe, to satisfy some standard of utility. In this sense, we (the putative neuronal mechanisms of our cognitive brains) do not discover the objects of common discourse; we create them for our individual and social purposes. This conclusion is consistent with Putnam's (1988) suggestion in support of the philosophical stance of internal realism, that "truth does not transcend use."

Given these biological constraints, it is clear that meaning and definition, expressed and understood through the medium of a common language within a community of individuals, can be no better than approximate and occasionally significantly divergent. This is true because among different individuals, neuronal tokens that are linked to identical lexical items on the output side are not likely to be linked to identical object representations (images) in long-term memory or to be embedded in semantic networks with identical associative (synaptic) structures. Furthermore, significant internal images may represent things that do not and never have existed in the real world. In such cases, referential content can be communicated only by extended description or by an analogical externalization like a diagram or some other physical artifact. If purposeful communication is a goal, there is no biological apparatus for ensuring a commonality of understanding. Our plans for communicative expression are shaped by the pragmatics of social convention and our perceptions of the practical consequences of individualized efforts.

Creativity

The capability for invention, trivial and great, is arguably the most consequential characteristic that distinguishes humans from all other creatures. Our cognitive brain is especially endowed with neuronal mechanisms that can model within their biological structures all conceivable worlds, as well as the world we directly perceive or know to exist. External expressions of an unbounded diversity of brain-created models constitute the arts and sciences and all the artifacts and enterprises of human society.

It is primarily the representational flexibility of the retinoid system coupled with the capacity of the synaptic matrices (including the semantic networks) for learning, long-term memory, and imaging that make creative modeling possible. An important basic function

of these putative brain structures is the inventive synthesis of specialized retinoid patterns of autaptic cell activity that can be transferred to long-term memory in the detection matrix for use as perceptual tools. These gaugelike tools (templates) can be used for quick and rough assessment of significant geometric properties in the visual environment. Within a more general perspective, the synthesis of novel and useful neuronal patterns that would be committed to long-term memory in the individual and shared cumulatively in the larger society would represent the evolutionary development of biological, intraorganismic cognitive tools in the same sense as the evolutionary development of material extraorganismic artifacts (Trehub 1977).

The Brain's "I" and States of Belief

I strongly suspect that the pragmatics of the cognitive brain model can help resolve current philosophical disputes about intentionality (Churchland 1984, 1986; Putnam 1988; Pylyshyn 1986), but I will not pursue this line directly because I have no clear sense of the norms of engagement on the issue. Rather, I will argue that certain properties of the theoretical model that I have already delineated provide sufficient biological grounds for generating and supporting a distinct internal token that warrants the lexical designation "I." This special token (the I-token), when joined in an internal stimulus complex with other active tokens, establishes a distinct internal state that warrants the lexical designation "belief."

I-tokens are a special subset of neurons characterized by having an excitatory input from the autaptic cells that constitute the self-locus—the source of constant excitation in the center of the self-locus retinoid. The discharge of an I-token can be thought of as a neuronal signal of oneself taken as an object. Because it is a labeled line, the I-token can be linked to the character "I" in the lexical assignment matrix and can serve as the subject in all kinds of subject-predicate propositions within the semantic network. Any recognized internal state or perceived event involving oneself can contribute to a store of learned predicates uniquely induced by activation of the I-token—for example, *I am sad* or *I travel a lot*. I-activated predicate tokens, in turn, can evoke sensory images by their chain of backward links to the imaging matrix in a level-1 synaptic matrix. The extended set of such neuronal associations can be taken as the biological substrate for one's sense of self.

In addition to signifying the center of the self-nexus, the I-token can mark a subset of neuronal propositions that have a very special status in the cognitive brain—those that have passed a test of per-

sonal validity. Consider the proposition "Kumquats are smaller than oranges." Now imagine two different individuals, each well acquainted with oranges, who read the sentence. One has never before seen, read, or heard of a kumquat; the other has seen the fruit and has learned that it is called a kumquat. The first person knows an assertion has been made that something called a "kumquat" is smaller than an orange, but there is no test that he or she can perform to establish that the predicate "smaller than an orange" is a true (sensory-world) property of the subject "kumquat." The second person can recall an image of a kumquat and actually compare its size with an imaged exemplar of an orange. When the outcome of this personal test establishes that a kumquat is indeed smaller than an orange, the proposition assumes the status of a truth within the cognitive system of the individual making the test.

I suggest that propositions of this kind are gated to a privileged semantic network that contains what can be appropriately called the system of personal belief. I also suggest that every proposition evoked from the network of belief is neuronally marked by the joint discharge of an I-token (I!). Thus:

((Birds) <fly>))I!

<People> <fly>

((People) <walk>))I!

<I> <can read Sanskrit>

((I) <can read English>))I!

<Dan> <needs a new jacket>

((Dan) <needs a new jacket>))I!

((I) <think Dan needs a new jacket>))I!

<I> <think Dan needs a new jacket>

<Dan> <says he needs a new jacket>

((Dan) <says he needs a new jacket>))I!

Neuronal propositions accompanied by the parallel discharge of an I-token are expressions or states of belief; those that are not accompanied by an activated I-token are subjunctive expressions or states. Notice that identical sentential propositions can differ in intentional stance, expressing a state of belief in one instance and a subjunctive state in another. Detailed consideration of this aspect of the cognitive brain would take us well beyond the intended scope of this book. It

is worth emphasizing, however, that our ongoing interpretation of the world and our formulations of plans of action are critically dependent on our prior beliefs. And these, in turn, depend on individually learned routines (not necessarily rational) as well as innate mechanisms for signaling the ostensive validity of propositions.

Conclusion

My proposed theoretical model of the cognitive brain consists of a number of putative neuronal mechanisms that together constitute the biophysical and structural basis for certain essential cognitive properties of the human brain. An interesting aspect of the component mechanisms, and a fact that supports their evolutionary likelihood, is that most of them are replications or thematic variations of either one of only two principal neuronal designs: the synaptic matrix or the retinoid. Current neurophysiological and neuroanatomical knowledge tends to confirm the biological plausibility of the hypothesized mechanisms. Computer simulation tests of the mechanisms demonstrate that they are competent in basic visual-cognitive and lexical-semantic domains. The extended integrated model of the cognitive brain is responsive to all of the ecologically relevant tasks listed in chapter 1 as reasonable tests of a human cognitive system. In addition, the model accurately predicts many classical illusions and perceptual anomalies as epiphenomena of its normal operation. Finally, a large body of diverse experimental and clinical findings can be explained within a coherent framework by the biophysical and structural properties of the model.

On these grounds, I believe that the neuronal mechanisms and systems that I have proposed represent a credible biophysical explanation of human cognition in the domains assayed. Although a number of important processes have been covered, many more remain to be explicated, although I expect that the mechanisms that have proved competent for the tasks considered in this book will, with some minor changes and elaborations, be equally effective for other modalities and tasks.

We are on the verge of a biological understanding of human cognition, but we have only a dim appreciation of the opportunities and challenges that such an understanding entails. As competing theories of cognition appear, we face the scientifically exciting task of deciding among them on the basis of competence, biological plausibility, and parsimony. This effort can illuminate and help clarify the daunting complexity of cognition and sharpen our sense of what it means to

be human. Reasonable models of the human brain can provide a physical account of the mechanisms and processes of cognition. Interactions among such models in simulated social and physical contexts can provide an account of the evolution of the various contents of cognition. It is the total specific content of cognition, the current physical state of specialized mechanisms in an individual brain shaped by encounters in a world both real and imagined, that constitutes a mind.