

Cortical Architecture and Self-Referential Control for Brain-Like Computation

Edgar Körner, Gen Matsumoto

2002

Preprint:

This is an accepted article published in IEEE Engineering in Medicine and Biology. The final authenticated version is available online at:
[https://doi.org/\[DOI not available\]](https://doi.org/[DOI not available])

Edgar Körner¹ and Gen Matsumoto²

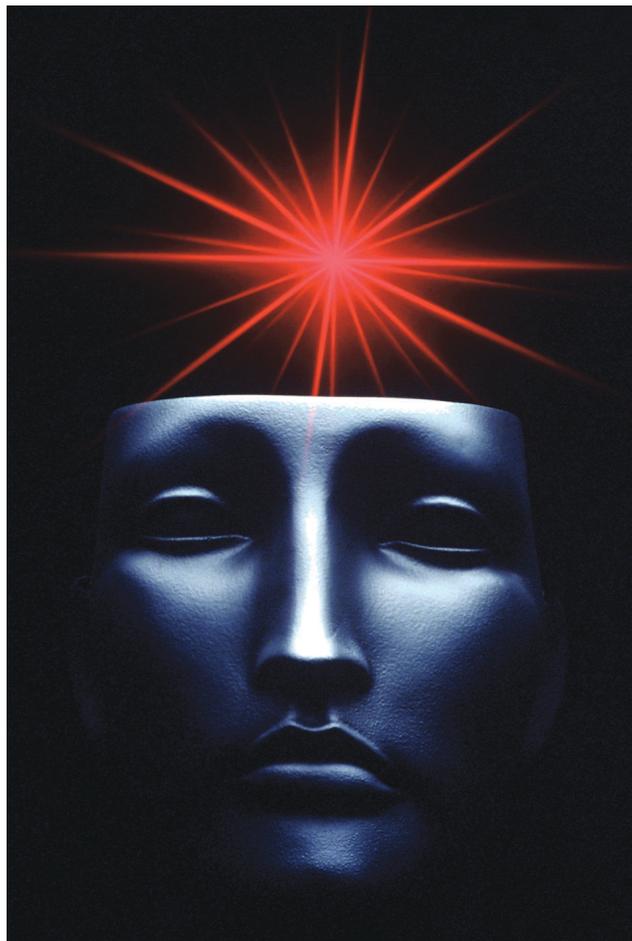
¹Future Technology Research Division, Honda R&D Europe

²Brainway Group, RIKEN Brain Science Institute

Cortical Architecture and Self-Referential Control for Brain-Like Computation

A New Approach to Understanding How the Brain Organizes Computation

Considerable progress in experimental neuroscience has not yet resulted in an enhanced understanding of how the brain organizes computation. In this article the conceptual problems behind most of the current brain models are discussed, and constraints for an alternative approach are defined. We propose an approach that may provide the framework of orientation for the development of brain-like computation. We argue that self-referential control architecture is an essential functional characteristic for an autonomous intelligent system. The constraints for self-organization of a brain-like neural system are also discussed. We shall demonstrate that the phylogenetically correct order of development of control in the cortex should be a *top-down* process instead of the widely accepted bottom-up scheme. Only with these constraints is an autonomous bootstrapping of knowledge representation feasible. However, before the system's self-referential control can start to bootstrap a consistent knowledge representation, at least some a



©DIGITAL STOCK

priori knowledge must be present to initiate the self-organization process. We shall propose a mechanism that might serve to initiate the self-organization of knowledge representation whereby a set of intrinsic values (as expressed in the emotional system)

fulfills the requirement of a priori knowledge. We will also outline a cortical architecture that can support the self-referential control of learning and recall for the case of visual recognition.

Need for a Novel Approach to Brain-Like Systems Impasse in Research of Brain-Like Computing

In recent years, neuroscience has made major advances regarding both investigation methodology and insights into local mechanisms of sensory information processing in the brain. However, processes such as exactly how the brain manages to recognize a familiar face of a person, or how a person manages to navigate effortlessly through a busy street, continue to elude us. Obviously, our current computational models of the brain have not yet properly addressed the essential features of neural processing. We have learned, in dealing with artificial neural networks, that the architecture of a network is a graphical

Many current artificial neural network models and brain dynamics theories seem more motivated by attention to mathematics and simulation techniques than to a genuine interest in understanding the brain function.

notation of the class of algorithms implemented in that particular network, and that the current formal neural networks represent approximation algorithms well known in engineering for years. From both of these insights, a straightforward explanation is suggested for the failure of neural networks to achieve the level of sophistication in performing tasks that are easy for an animal in everyday life but are hopelessly tricky to implement with current computer technology. When accommodating current formal neural networks to represent certain phenomena of neural structures, we actually use the kind of algorithms that are optimized for the von Neumann type of computation as a nonparametric model for mimicking the observed phenomena. However, if we indeed want to get a deeper grasp on how the neural systems manage to accomplish those marvelous and robust performances, we need to take a closer look at both the physiology and morphology of the real human brain.

If we assume that neurobiological processes have no other properties than that which current computing technology al-

ready replicates (e.g., assembling, matching, and storing of signal or symbol constructs) and we limit the brain computing models to what current computers are capable of doing, we will be prevented from gaining the insights into a) how the brain may avoid the inflexibility of rule-bound mechanisms, b) how it may escape the pitfall of combinatorial search, and c) how it solves the problem of smooth transition between signal and symbol processing, efficient real-time coordination, etc. Certainly, there is a major difference in organization. The wiring of the brain circuits is highly variable, and the representations change with the passage of time. Behavior—hence, also the computation that generates just that behavior—in neural systems seems, to some extent, self-generated in loops. Sensory input leads to brain activities that result in behavioral responses, which lead to further sensation. The latter, in turn, leads to further brain activities, and the cycle goes on and on. A process comparable to software compilation does not implement such a rearrangement of neural system components to generate appropriate behavior. Each perceptual categorization and each sensory-motor coordination links the brain's "hardware" components together in new configurations, creating new architectures within the population of available physical elements, for further activation and recombination. Thus, to better understand the type of computation inherent in the brain, we have to contrast neural computation against the programmed devices that adapt within fixed, predesigned constraints, instead of explaining brain processing in terms of those systems.

Utilizing an Evolutionary Perspective to Model Brain-Like Systems

Implementing the essential characteristics of functional architecture of neural systems is a prerequisite for the construction of a machine that can rival the brain for its cognitive functions. Starting from regularities revealed by morphology appears to be a good idea if neural architecture reflects the algorithms implemented therein, as outlined above. However, a close reference to neurobiology is a necessary but not sufficient prerequisite. Neural systems must be investigated within the context of a purposeful organization of behaviors. Even a bird's wing would be a mystifying structure if we did not know its purpose. Defining a truly challenging problem of recognition or

behavior, and then attempting to understand how the system could solve the problem under the constraints, imposed by the functional architecture and the existing neural hardware, will be a more promising approach than defining the problem in terms of algorithmic structures known from current computer science. For a complex system like the brain, functional organization matters, and it cannot be ignored without rendering the model irrelevant.

In following that approach, how can we know what are the relevant control circuits for a particular function when we are confronted with the almost incomprehensible complexity of the brain architecture? Until now, we do not even understand the functional organization of a single neocortical minicolumn, which is an elementary subsystem of cortical processing. Already at that elementary level of organization of neural computation, there are many intertwined controls for diverse functions that we cannot decompose into meaningful components without having sufficient insight into some fundamental principles of how the brain organizes computation. Because of the great complexity in a biological system, extended chains of mathematical reasoning are less relevant in biology than in physics and engineering, unless it is corroborated at every pertinent step by experimental data. However, many current artificial neural network models and brain dynamics theories seem more motivated by attention to mathematics and simulation techniques than to a genuine interest in understanding the brain function. "Neurally inspired" formal neural networks (see Haykin [1] for an overview) and connectionist semantic models (e.g., see Hinton [2] and Shastri and Ajjanagadde [3]) evade a direct relation of its nodes to real neurons in terms of what they actually represent, and how they may be embedded in a purposeful architecture in order to enable the flexibility of processing—a desirable feature to implement in artificial systems. But we must confront such a question if we are to develop a realistic theory of the brain. Nonphysiological models may provide interesting metaphors to explain isolated phenomena, but they are unlikely to make a major impact on our understanding of how the brain works as a *whole*.

The solution may rest with understanding the development of the systems in a bottom-up fashion. Understanding the way a system developed from a simple reflex control architecture to a sophisticated intelligent system by adding an extension of the

control hierarchy, whenever the demand for a more advanced behavior arises, will dramatically reduce the complexity that confronts us when we examine the final architecture. It appears that the order of development of controls is critically important. In the ontogenetic development of brains, the phylogenetic development is repeated in a “fast-forward” mode. We argue that it is impossible to develop the complex hierarchy of diverse varieties of controls concurrently in a *single* step. When we model brain systems, the sequential order in which modules of control hierarchies are developed may be crucial for attaining the level of flexibility and robustness known in a biological system. Moreover, the well-orchestrated development from simple to complex control systems may be critically important for ensuring *stability* of those many intertwined controls. For instance, flapping wings and feathers turned out to be irrelevant toward the design of an airplane, but the *phylogenetically correct order* of implementation of the behavioral capabilities and its related controls is: first the airfoil principle, then dynamic balance, and last propulsion [4]. That means the *order* of acquiring various capabilities matters. We should not only focus on a final version of a sophisticated and fixed architecture, but also we must find out how the architecture evolved, as its behavioral capabilities emerge from a simple to a complex one.

Self-Referential Architecture and Control

Complex systems like the brain must have an enormous number of control mechanisms that are not directly related to the *content* of sensations, categorizations, etc., but rather to the *process of making* these sensations and categorizations possible. This provision ensures a unitary organization of the system, which enables its coordination to adapt to the changing environment, such as synchronizing, managing, and arbitrating among different functions. The brain has evolved under these constraints set by this type of the global control. Hence, without knowing at least some general characteristics of these control procedures, and the order in which they developed, we may fail to address the essential aspects of the system’s function.

How could we implement such capabilities in artificial neural systems? As outlined above, the basis of cognitive abilities is not the facts memorized in the course of sensory experience but rather the architecture and the self-referential

control that dictates the brain to make these representations. The architecture of the brain should not just reflect the knowledge stored there, but it should reflect the control that provided the constraints for its acquisition. This is the unique feature of the brain’s organization, as compared to a conventional computer. Biological systems represent the world (including themselves) in a way that makes possible the optimal control of behaviors by means of *prediction*. The iterative tuning of this prediction successively minimizes prediction errors. “Self-reference” is a crucial control principle to enable this type of cognitive ability: any modification of the sensory input states produces a prediction error—since the previous prediction was based on previous sensory inputs—which the control (“self”) must subsequently compensate with a new and modified prediction. If the prediction error cannot be sufficiently compensated by one selected from the existing repertoire of predictions (representing the “already-acquired” knowledge), the prediction error and the sites in the representational framework, where those errors occur, define what is to be learned and where to implement.

Instead of the entire situation that caused the prediction error, only the difference between the prediction and the actual situation needs to be learned. In this way, the system always makes a self-reference when it functions in the environment: the inability to predict the sensory input by the existing knowledge defines what new knowledge needs to be acquired and where it is to be learned. Hence, only these “difference” aspects are added to the existing knowledge representation. In this way, new knowledge is always integrated into the relational framework of existing knowledge, for uses in tuning up future predictions.

The simplistic model of sensory processing, reduced to a parallel recall of stored pattern knowledge triggered by the sensory input, cannot describe the comparison of sensory input with prediction. This purely forward processing approach is related to the “feature detector” philosophy in experimental and cognitive neurobiology [5], [6], where simple and mostly artificial stimuli are presented to a sensory input and the response characteristics of a neuron are measured. However, usually a natural input does not consist of an isolated single object but of complex scenery composed of many objects which vary considerably over time even if the basic scenery does not change.

Experimental evidence suggests that, at least in awake subjects, there exists an internal representation of the external world.

Solving the recognition problem by purely forward processing in the pattern-matching paradigm is impossible for the brain for at least two reasons: First, the brain cannot have a parallel (pattern) representation for each and every possible variation of such complex scenery in which the living being has to behave regularly. Second, even if it had indeed the tremendous memory capacity to do that, it would be so frequently locked in the impasse of combinatorial explosion that a fast reaction (the ultimate goal of recognition) would be impossible. Continuous self-reference as proposed above could serve to overcome those critical limitations.

Experimental evidence suggests that, at least in awake subjects, there exists an internal representation of the external world [7]. That means the brain is not a machine that is silent until a suddenly appearing stimulus triggers some short activation to respond to that input, but it is always actively keeping its internal understanding of the current situation to evaluate its sensory inputs from that point of view. There is also evidence that recognition in the visual cortex is a *bi-directional* process, in which partial results of bottom-up processes trigger higher-level object representations, which then in turn guide the segmentation process to decompose the scene into meaningful objects [8]. Consequently, this segmentation must be a sequential process of iterative refinement of recognition by top-down prediction, starting with the most salient feature of a scene, and subsequently adding more and more details, from a coarse to a finer and finer description of the input. The most salient feature (be it a general framework of the scenery or a specifi-

Our hypothesis implies that emotion is not a mere “side-effect” of brain processing but is actually the crucial basic knowledge for relating the sensory input patterns to an appropriate behavioral prototype.

cally attended object in a certain context) that is recalled first in the bottom-up filtering process constitutes an *initial hypothesis*. This initial hypothesis is fed back as a top-down prediction and thus controls the subsequent refinement process [9]. The decomposition is driven by comparison of sensory input with the brain's prediction until there is no residual part of the input that is not yet accounted for, or until there is no internally existing representation that can account for this residue. This not-yet-known part is interpreted as a *new* attribute in the context of the assembled prediction and added as a new item to the sequence of decomposed semantic entities [10]. Later, this attribute may be decomposed even more if there is a behavioral need to do so. Thus, the detail is added *on demand*. There are two major characteristics of this “self-referential control” of the cooperative bi-directional recognition process, which constitute prerequisites for autonomous behavior:

- The decomposition of any input into existing representational primitives enforces both efficient and consistent knowledge representations.
- The activation of the initial hypothesis links the current sensory input to a behavioral quality, hence giving semantic meaning to that input. Consequently, its control of the subsequent decomposition process enables the utilization of existing knowledge (which is equivalent with the interpretation of the input in terms of the system's semantics) to allow for prediction in an unknown environment.

Hence, the system can even react to unknown events by trying to find the best analogy (“nearest,” though not perfect, match) existing in the repertoire memorized from its experience. The recall of already-acquired knowledge in response to a sensory input provides a value-based behavioral quality to assess the semantic quality of the sensory situation. Furthermore, by using the recall as a prediction to be compared with the forward description, the brain has the capability to define the degree of fit between the *internal world model* and the *sensory reality*. That is why the brain can learn without a teacher, since it can apply a *gradient strategy* for improving its predictions by successive approximation starting from an “initial hypothesis” [11].

The next questions we will address are the definition of a “self” in a brain-like system and how to provide the higher levels of the systems control hierarchy with the minimum required knowledge to start the self-organization of the knowledge representation.

Brain-Like Neural Systems that Self-Organize Their Knowledge Representation

Defining a “Self”: Setting a Semantic Metric by an Intrinsic Value System

The cortex is generally assumed to be the site of the highest hierarchy in the mammalian brain for decision making, because it is known to be the site of the most elaborate processing and decision making. However, evaluation of the information from the environment and the selection of an appropriate behavioral response is well managed by phylogenetically older brains, even if it was not as elaborate. We propose that the phylogenetically older structures, which selected appropriate behaviors in “lower” animals, are still in command for brains with cortex in higher animals, setting the constraints for the cortex regarding what basic behavioral meaning it should extract from the sensory input, among the

many possible alternatives for an interpretation. We propose that these phylogenetically old structures are still at the *top* of the processing hierarchy. Hence, we will refer to these structures as higher centers, contrary to the convention in neuroscience. They assume command at least for the start-up phase of the development of an internal knowledge representation. This ensures an elaborate relation to be established from the sensory input to the behavioral response. They also take command in time-critical behavioral situations where a reaction has to be selected promptly. The evaluation of the sensory input with respect to the system's intention must occur in direct reference to the *limbic system*. The limbic system belongs to the phylogenetically oldest parts of the brain and can be considered as the “value system.” The limbic system deals with emotions and performs a coarse judgment of the state of the individual organism relative to its environment: it evaluates a state as “good” or “bad” for survival, or simply whether the state is palatable to the present state of desire of the animal. Only recently has it been understood that emotion provides the coarse semantic metric for the mapping of sensory input to fundamental types of behaviors [3]. Philosophically, this value system, as compared to other parts of the brain, may be most closely related to the concept of “self.”

The brain acts in accordance with the changes of its environment, based on the evaluation of the current situation from its “self”-interest, such as protecting its own existence, its own well-being, etc. Probably it is this “being a self,” rather than a merely reactive automaton, that is at the root of flexibility in information processing and behavioral responses. The system needs not be told of any categorization. It is always interpreting the objective syntactic sensory information from the viewpoint of its subjective internal description of the external world, including itself.

The internal representation of the outside world is selectively defined by the value system in accordance with the behavioral needs of the animal that is the host of the “self.” This permanent *subjective* re-description of *objective* sensory information into a *consistent* internal representation of the world seems to be a key process to enable robustness and flexibility, which is evident in solving the symbol grounding problem and in enabling rapid coordination of available resources for smart real-time behavior. Enforcing and keeping consis-

tency of this internal representation requires powerful control of self-reference, since any new information has to be evaluated in terms of the previously acquired knowledge, and must then be integrated into the existing relational architecture of the subjective knowledge representation.

The Neocortex Developed Top-Down: Learning Is Controlled by Top-Level Semantic Metrics

The self-referential architecture and controls must exist from the beginning of the self-organization process. Hence, they should be encoded genetically. However, since all brains with cortex show the capabilities of flexible behaviors, the key elements of the flexible behaviors should have been expressed in the phylogenetically oldest cortical structures already.

The failure of top-down designed systems of early AI to ensure flexibility and robustness when confronted with real-world problems in recognition and control bred the bias toward purely bottom-up schemes of self-organization of knowledge representation, as exemplified by Marr's approach in vision [12], or the subsumption architecture in robotics proposed by Brooks [13]. However, when the development started at the lowest sensory and behavioral level and proceeded to increasingly more complex sensory and behavioral repertoires, the direction of the self-organization of the architecture must be defined by trial and error, according to this bottom-up scheme. Neither statistically based learning nor evolutionary learning (e.g. genetic algorithms or evolutionary optimization rules) are sufficiently effective for such a purpose, especially not for handling more complex problems in sensory categorization. To accomplish more than simple reactive behaviors, the designer must specify the jobs and the necessary conditions when adding a higher level to the systems hierarchy. The philosophy behind these conventional approaches is the assumption that the brain has developed bottom-up and that the emerging higher levels of the system's architecture subsequently *enslaved* the lower levels. However, our present understanding is opposite to this conventional view. We argue that the phylogenetic development of brains proceeded indeed in a bottom-up fashion with respect to evolution from low to increasingly higher competence for the interaction of the living being with its environment, but in a top-down fashion for the evolutionary refinement of sensory

analysis and behavioral synthesis within a certain competence level.

Recent insights into phylogenetic development of brains with cortex suggest that there is actually a *top-down* process [14]. The phylogenetic development of the neocortex started from two prime moieties: the paleocortical moiety (temporal pole) tied to holistic sensory analysis, and the archicortical moiety (hippocampal cortex) tied to generating behavior. From these moieties, all other cortical areas developed, step by step, with interconnections between areas of the same level of the emerging hierarchy [15]. Hence, both for sensory analysis and behavior generation, the "highest" level of the hierarchy dealing with the most general evaluation of the sensory situation from the subject's point of view has been in place first. They controlled the subsequent correlated development of lower level of analysis and behavior generation. The phylogenetically youngest levels (last in development) are the primary sensory and motor areas. Only the oldest part of the cortex (the origin of its phylogenetic development) is reciprocally connected to amygdala (AM). AM is an old subcortical part of the limbic system related to emotion, and it has been shown to trigger elementary behavioral categories (e.g., fear conditioning, arousal, and escape behaviors). Based on this line of experimental evidence we propose that top-down development of the brain proceeded under the control of AM and related structures. AM and related structures constitute a *value system* for the subject (the living organism). We conclude that the top-down control of the primate neocortex should still be dominant in defining the direction for understanding a sensory situation, thus setting the stage for the interpretation of the sensory input by the lower-level systems (Figure 1).

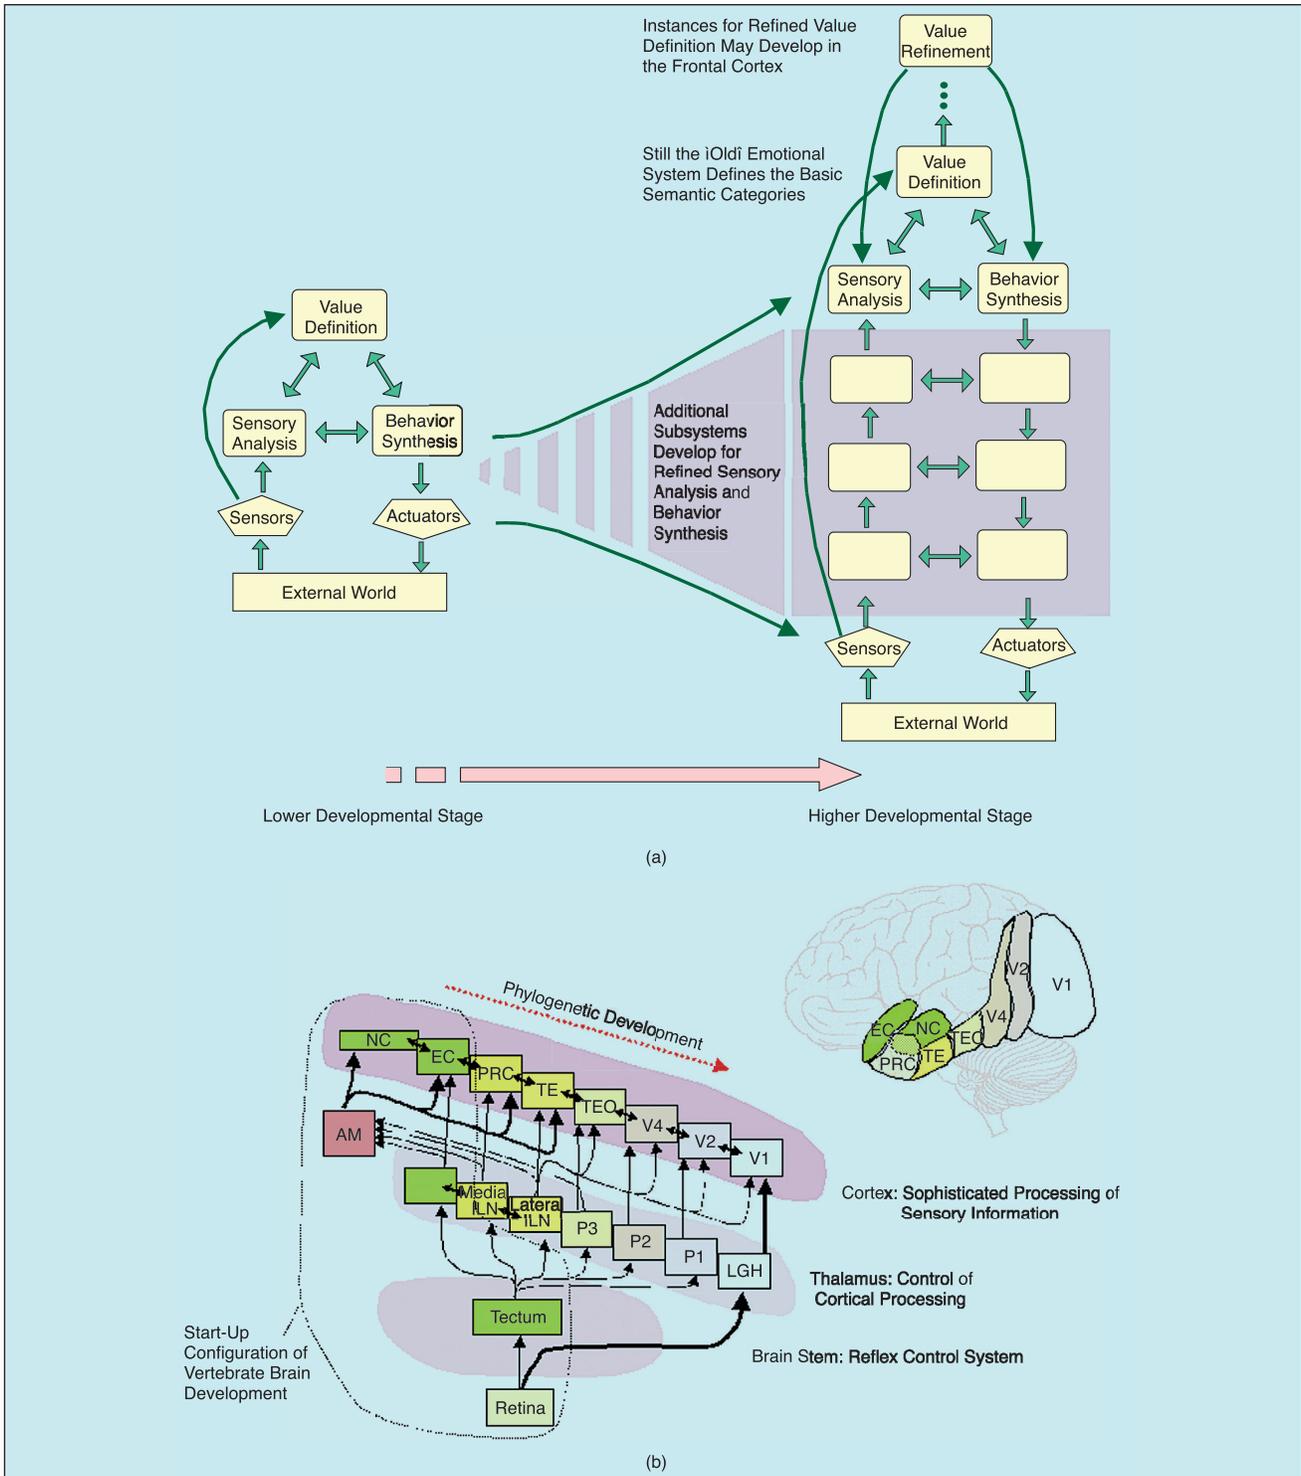
Next, the bootstrapping of a relational knowledge representation enforced by such a self-referential control could be plausibly explained as being guided by the previously acquired knowledge that serves as an initial hypothesis, even if it is not very sophisticated or accurate [10]. However, some a priori knowledge must be made available to the system to "initialize" the process. Since this a priori knowledge is a prerequisite for integrating sensory information into a consistent internal representation, it must be present even before any sensory experience can shape the brain. We conclude that not only the control architecture but also some "initial knowledge" must be genetically encoded.

Example: Scheme for Self-Organization of Knowledge Representation in Vision

With the development of systems serving an increasingly detailed analysis of sensory inputs, the sensory input to the cortex may have been moved downward together with the emerging lower hierarchic levels to the current primary sensory areas. However, the "old" input to the phylogenetically oldest subsystems of sensory analysis is still present as a coarse but direct input via AM, CS (colliculus superior which is a homologue of tectum), and thalamic pulvinar nucleus [Figs. 1(b) and 2]. In the case of the well-investigated ventral visual pathway of the cat, there is a direct input from sensory thalamus to infero-temporal (IT) cortex and all lower levels of visual filtering, although the main stream of sensory input enters the visual cortex at the primary visual cortex (V1). Note that V1 is the phylogenetically youngest area that represents the lowest level in the hierarchy of cortical visual processing. To initiate the self-organization process of sensory categorization and goal-directed behavior generation, some a priori knowledge has to be available within the system. Such a basic set of a priori rules might be stored in AM [14].

There is a wealth of experimental evidence, suggesting that AM neurons code a set of behaviorally important coarse sensory situations [16]-[18]. The shortcut sensory input to the highest level of unimodal sensory processing (inferotemporal cortex, IT) is formed before the developing neocortical afferent input, via the primary sensory areas, and reaches the same target [19]. Coarse representations in IT are already established in 10-day old infant monkeys—a time when the developing afferent bottom-up connections in the neocortical filtering hierarchy have not yet reached IT. These representations might be generated by subcortical inputs and are not subject to modification of subsequent sensory experience when the *neocortical afferent pathway* to IT is fully developed [20]. The tentative source for *subcortical input* is the direct input from thalamus via AM. Projections from AM to the phylogenetically oldest part of IT are reciprocal and topographically organized [21]. This, however, does not apply to connections from AM to other sensory areas [22].

Based on these findings, we propose that AM may serve as a "teacher" (semantic pointer or trend setter) for establishing representations of sensory environment in



1. Schematic drawing of phylogenetic top-down development of the cortex. (a) The phylogenetic development of the cerebral cortex seems to have proceeded in a top-down fashion, starting from the “archicortex” that controlled a behavioral repertoire of limited complexity, and a sensory analysis of limited resolution that was sufficient for that behavior. At the left side of the figure, the start-up architecture is shown, where the emotional system already defines the semantic interpretation of sensory input based on the genetically imprinted initial knowledge (see text). During the phylogenetic development, the necessary refinement of the behavioral repertoire is consequently accompanied by a similar refinement of sensory analysis. However, these expansions of processing architecture are not made at the top of the hierarchy. Additional levels of processing are inserted between the sensor/actuator level and the “archicortex” creating a multilevel hierarchy that serves increasingly refined processing while the cortical and subcortical parts of the limbic system, the “archicortex,” are still in command, as shown in the right hand part. The horizontal (pink) arrow indicates the direction of phylogenetic development. (b) The detail of phylogenetic development of the cortex, which subserves vision, is shown. The inset (right upper corner) shows the anatomic localization of various parts.

IT, which then could be refined and diversified by the rich details delivered later by forward processing along the cortical hierarchy. These representations can actually guide the afferent (bottom-up) processing by sending feedback to lower levels of sensory filtering, since the established coarse representation at IT is always activated first, even before the afferent wave of filtered sensory details arrives at IT via the cortical hierarchy. The experimental evidence cited above also implies that a basic set of such sensory activation patterns linked to prototypical behavior should be genetically imprinted in AM. We conclude that these rules could serve to start building a knowledge representation in exploring the environment, defining the coarse constraints within which the system can self-organize in accordance with its sensory experience.

AM seems to be the site that determines the semantic metric at IT for cortically preprocessed sensory patterns. The term “semantic metric” is used here to describe the relation between a certain situation of the environment and the activation of the appropriate behavior. AM as a key element of the emotional system is proposed to define the fundamental categorization of sen-

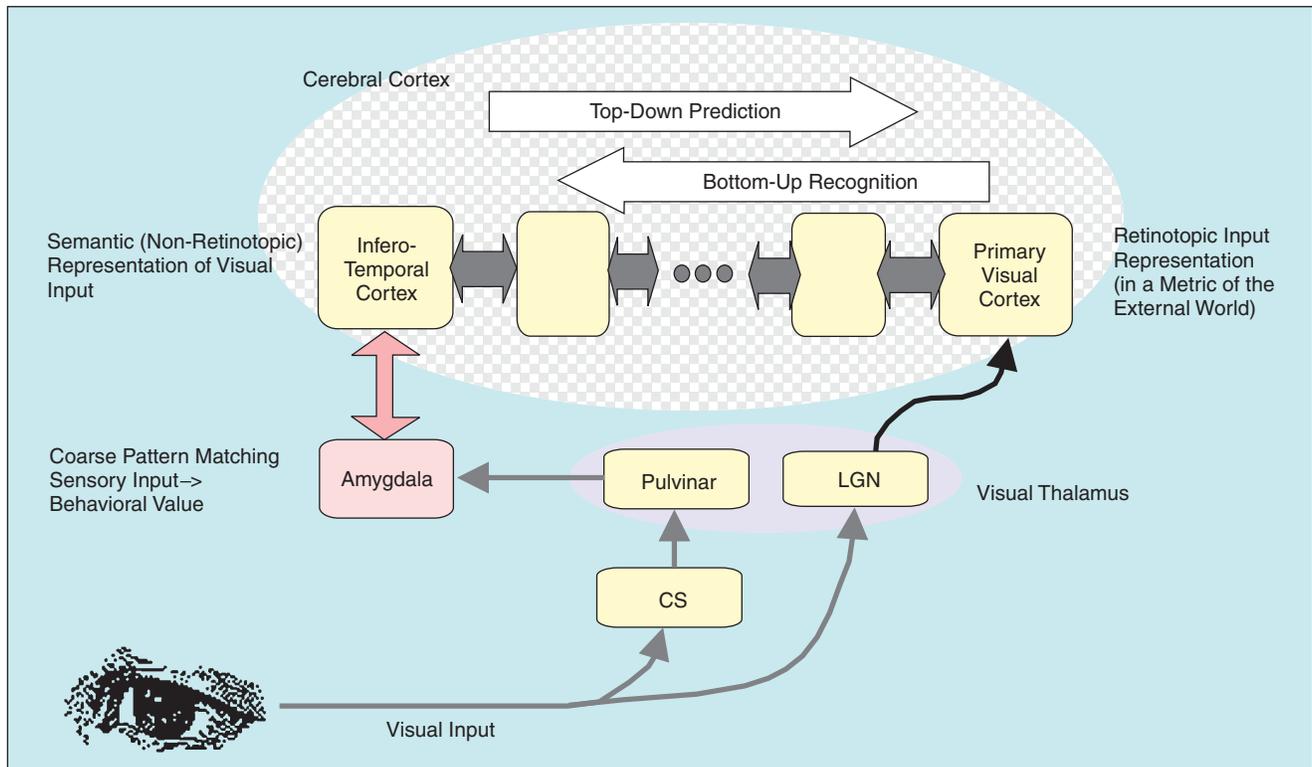
sory situations (good or dangerous, pleasant or annoying, etc.), providing a coarse semantic metric for interpretation of sensory input. Then, ongoing experience of the animal can serve to generate a more and more elaborate and sophisticated hierarchy of both sensory representations and behavioral repertoires.

Moreover, since input to IT via the shortcut pathway, thalamus \rightarrow AM \rightarrow IT, is much faster than via the refined cortical filtering, this shortcut may serve not only as a semantic pointer (“trend setter”) for learning. In any recognition process, this input via a subcortical shortcut is manifested in the highest levels of cortical sensory processing (which are the phylogenetically oldest structures, such as the hippocampus and entorhinal cortex) much earlier than the detailed filtering upward the neocortical main pathway arrives at the same site (e.g., see [23]-[25]). Hence, it may also pre-activate a coarse representation at IT as a hypothesis on the sensory input. Based on the evidence furnished by LeDoux [26], we further propose that this activated hypothesis may have already defined a semantic evaluation of the sensory input and set the semantic constraints for the

detailed processing of the identical sensory signals which arrive somewhat later at the neocortex. The genetically imprinted a priori knowledge in the limbic system may serve to start the semantic categorization. However, the emerging refinement of sensory analysis and behavioral synthesis requires a similar refined value system that may develop in, probably, the frontal cortex and operate on behalf, but still under the supervisory control, of the limbic system.

Rapid Learning in Complex Hierarchies: “Now-Print” Command by Amygdala

Having a teaching input from AM that defines the site at the top of the hierarchy of sensory representation, where the input pattern arriving at the lowest level has to be connected to, is a proper prerequisite for self-organization of knowledge representation. However, sensory situations are highly variable and do not repeat frequently enough to facilitate statistical learning. For some sensory situations the subject has no luxury of having a second chance to recognize it (danger, etc.) and should not need it either. Learning to represent such brief



2. Setting a semantic metric for knowledge representation. The shortcut from visual sensory input via AM activates the highest level of unimodal visual sensory processing, the inferotemporal cortex (IT). At IT, visual objects are already represented (both phylogenetically and ontogenetically) before the developing cortical hierarchy of elaborate filtering provides input to the IT [20]. LGN stands for lateral geniculate nucleus. CS stands for colliculus superior.

manifestations and transient states in a heterarchy is not an easy problem.

We propose the following mechanism as a possible scheme for one-shot learning in the brain [14]. In the case of strong emotional activation (arousal), AM trig-

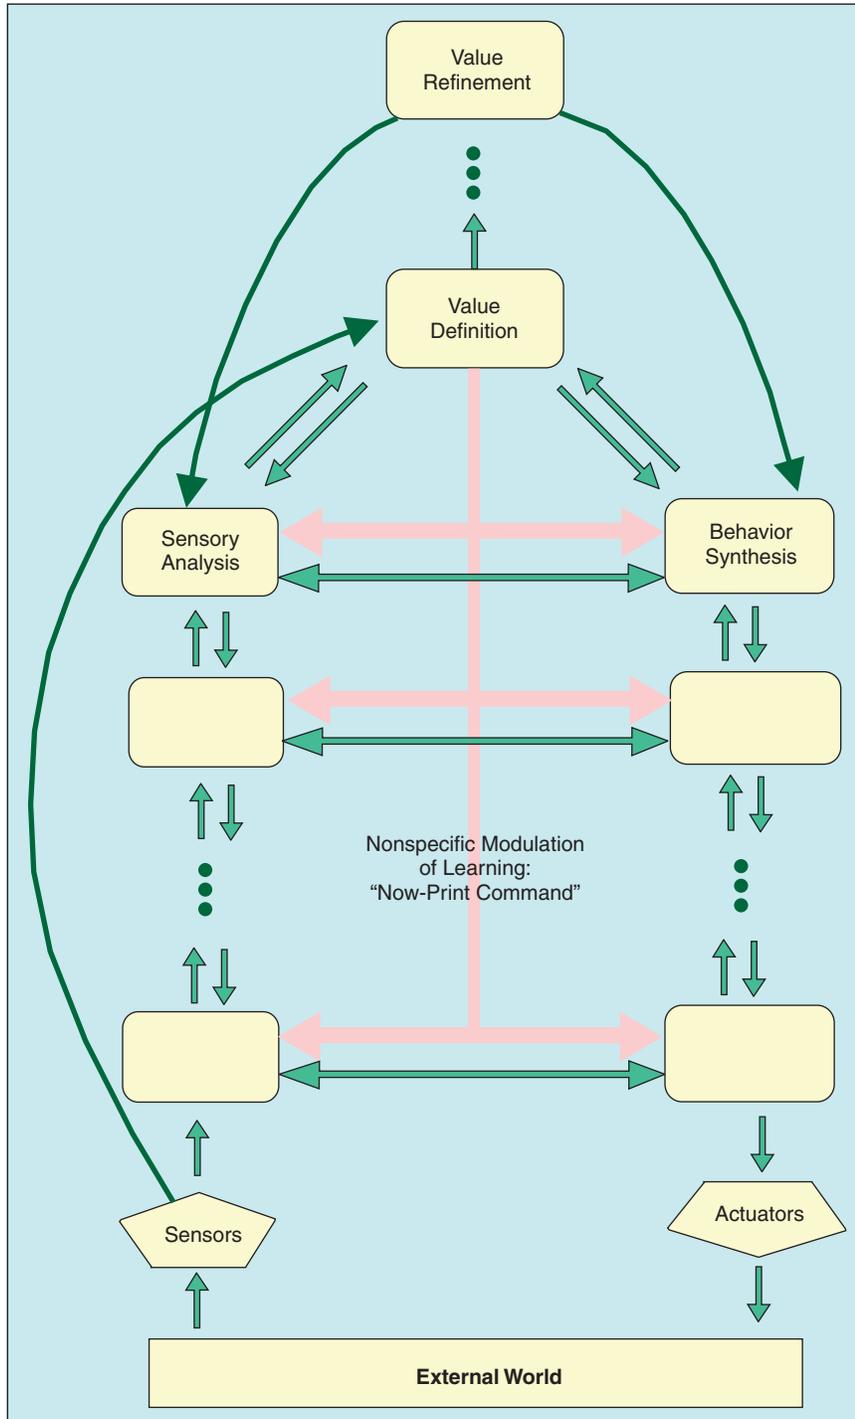
gers a nonspecific activation and supply of transmitters that support the formation of long-term potentiation (LTP) in the cortical visual and auditory areas [26], [27], [28]. The part of AM, which is topographically connected to the oldest part of

IT, may define where a certain sensory situation has to be represented within the internal metric of semantic categorization. The other nuclei of AM may serve to distribute a “now-print” command for one-shot learning nonspecifically into all cortical areas by distributing a transmitter that supports learning at synapses. These nuclei of AM release that transmitter in response to an arousal state. Arousal can be intentionally triggered by the system itself because the subject wants to learn the situation and pays attention to it. If the subject failed to respond appropriately to a sensory input since its internal representation of the situation has not yet been proper enough to elicit the required behavior, then disappointment or fear triggers the arousal (Figure 3).

Existence of a Value System Is a Prerequisite for Any Self-Organization

Our hypothesis implies that emotion is not a mere “side-effect” of brain processing but is actually the crucial basic knowledge for relating the sensory input patterns to an appropriate behavioral prototype. We propose that emotion serves to determine the coarse semantic metric for the evaluation of sensory situations and to support its refinement by sensory experience. We conclude that for any development of an autonomous system, regardless of what class of performance is selected, an initial value system has to be designed that sets the coarse potential characteristics of behavior within which the system can self-organize. Refined sensory analysis should develop in accordance to the need for refinement of behavior. As mentioned above, self-organization of sensory categorization cannot be understood in terms of a bottom-up philosophy. There are two key requirements for creating artificial systems that possess the kind of flexibility and robustness we cherish in relation to the performance of the brain of a living organism: There must exist

- a dominating top-down control for relating sensory events to meaningful behavior, and
- a “supervisory control” by a value system (like the emotional system) for coarsely setting the semantic metric and supervising its refinement with ongoing experience to maintain the consistency.



3. One-shot learning by nonspecific top-down reward control. We assume that a “now-print” command is issued from the amygdala to all subsystems of the heterarchy of sensory analysis and behavioral synthesis. The amygdala also activates the nucleus basalis Meynert, which gives a diffuse supply of the transmitter acetylcholine to all neocortical areas, as indicated by pink arrows. Acetylcholine is known to support the LTP formation not only in the hippocampus but also in the sensory cortex [28].

Robust and Flexible Recognition: Analysis-by-Synthesis

Predictions Constitute Dynamic Constraints for Top-Down Control of Interpretation of Sensory Input

A top-down control that guides the refinement of the systems architecture during phylogenetic development, as shown in the previous section, should also dominate the subsequent recall and learning of facts within the same architecture.

The recognition of complex sensory situations must be an active process, involving both bottom-up and top-down processing. The first stage in visual recognition seems to be a coarse holistic classification of the input [29], which is subsequently supplemented with details [30]. Accumulated experimental evidence from psychophysics supports the assumption that conscious experience involves an *active modeling* rather than a *passive perception* of incoming information [31]. A self-referential system will respond to sensory input by trying to find patterns and regularities that are consistent with the previously acquired internal representations, and to assemble these known aspects into the largest possible contextual framework. Those aspects that are already recognized should be separated from the input to facilitate the detection of new regularities [32], [33]. For this purpose, the neocortex has to construct a “working model” of the environment for comparison with the actual sensory input. Then it has to remove those parts that correspond to known items and finally analyze the residue. Therefore, any recognition must involve some kind of generation of internal hypothesis and, subsequently, iteratively verifying it by eliminating the previously known aspects from the sensory input until the residue vanishes or until it reaches the minimum, as limited by the individual’s ability. We conclude that interpretation of sensory input seems to be not merely the activation of a higher-level internal representation. Sensory interpretation is a process of re-creation of the internal representational architecture that relates best to the pattern of sensory input. This internal representation then serves as emerging dynamic constraints for further refinement of the recognition process (Figure 4).

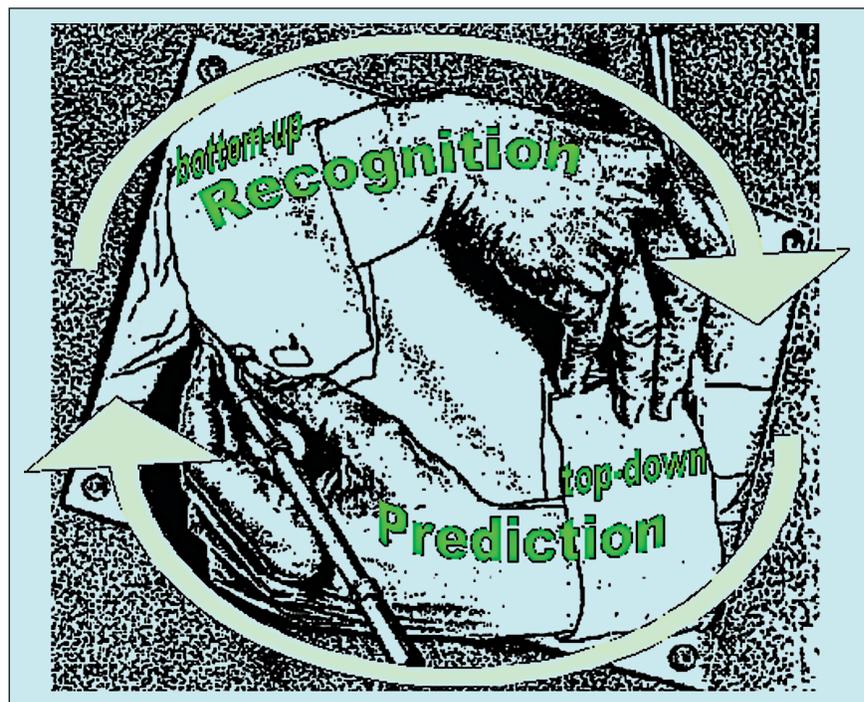
In this section, we only consider the case of visual image recognition, since the visual system is by far the best investigated neural system of the cortex. It is established

[34] that a proper global hypothesis on the sensory input will speed up an image interpretation process and will avoid the problem of combinatorial explosion, which is inherent to the bottom-up local filtering approach. However, to generate an initial hypothesis in the case of “real-world” problems like recognition of complex scenes, the prerequisite seems to be that one has already grasped the essential aspect of the scene—what is actually the outcome of the purpose of recognition itself. In other words, for generating a proper holistic hypothesis at the top of a hierarchic recognition system to control the processing at lower levels, one should have already solved the most difficult part of the recognition problem, namely the coarse understanding of the semantic meaning of the sensory input.

According to experimental evidence, we argue that two separate mechanisms serve to generate an initial hypothesis. The first mechanism is related to the shortcut input of raw visual data to the higher hierarchic levels of visual representation in the ventral pathway via AM,

which may activate a coarse understanding of the behavioral context of the expected information (see Figure 2). The second mechanism is a rapid forward transmission of most reliable decisions of local recognition units of the major processing stream of sensory information that activates a specific initial hypothesis at any level of the cortical visual processing hierarchy. The feedback of this initial hypothesis, which serves as prediction to the lower levels, sets the dynamic constraints for the subsequent refinement of this holistic but coarse hypothesis [9].

Neurons in AM show visual stimulus-selective responses, but they have a stronger tendency to respond to a specific category of stimuli. Some respond only to a limited category of stimuli, such as an angry face of a human, a particular person, a certain type of food, a threatening face of a monkey, etc. [18]. By the topographic reentrant connectivity to IT, AM could set a pre-activation according to a coarse semantic metric. Then, the most salient aspects of the input that are transmitted most rapidly along the neocortical filtering hi-



4. A sketch on neocortical computation. The famous Escher figure eloquently depicts the bidirectional characteristics of neocortical processing. We see what we expect, and we understand what we can predict. The recognition process is started by a rapid forward categorization that activates an initial hypothesis at each level of the processing hierarchy, which provides a top-down predictive feedback to verify this hypothesis. This process is assumed to mediate essential aspects of selective attention. Sensory description in the forward process and predictive feedback refine each other and create a consistent description of the input in terms of already acquired knowledge. Categorization can be done in feedforward processing; for discrimination and in case of complex segmentation problems, bidirectional processing is required.

Without proper
 “firmware
 organization,” artificial
 neural systems may
 reasonably perform a
 certain task, but they
 will never achieve the
 robustness and
 flexibility that
 we cherish in the
 “biological original.”

erarchy may specify a more definite but still coarse holistic description within an interval of about 100 ms at IT cortex in a forward processing [35]-[37]. The rapid activation of a holistic initial hypothesis on the raw input data at the top level of the visual pattern representation can impose *content-dependent* constraints on the af-

ferent input by *content-specific* feedback modulation to the lower levels of the neocortical hierarchy. In case of a simple feature or object, the top level of a holistic representation may be the IT, while the perirhinal cortex or, even higher, the entorhinal cortex may be the adequate location to encode a complex scene. We assume that the refinement of recognition at the lower level is only possible if a sufficiently entrained hypothesis at the respective higher level feedbacks a definite support [9]. This top-down prediction enables the activation of related details of the description, limits the range of search, and prevents combinatorial explosion of having to search all possible alternatives in the decision processes.

This assumption is supported by the fact that object representations in IT cortex are activated within 5 ms after the afferent information arrives and do not change subsequently [36]. In contrast, at V1 and extrastriate visual areas, the ensemble of activated feature detectors may vary within a 200 ms window of the sustained cortical activation after arrival of the afferent information [38], [39]. We propose that this 200 ms response duration reflects the refinement of the internal hypothesis in a process of hypothetical reasoning. The subsequent activation of the less salient parts of the input description is biased under the feedback from the more global evaluation at the higher level.

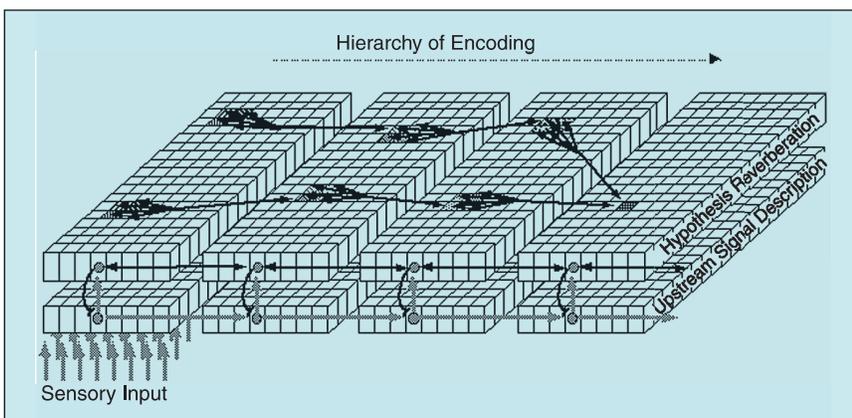
Generation and Verification of Hypotheses Require Complex Elementary Processing Nodes

In order to implement the self-referential control for recall and learning, the predic-

tion, generated by the activated hypothesis, must be compared to the forward input description (see above). Since the system must decide exactly where and to what extent the knowledge has to be updated in order to enable next time a correct prediction of the input in question, the comparison is to be made at any processing node in the systems architecture. This difference must be defined both explicitly and promptly. Usually, there is no chance to keep the input constant long enough to perform a kind of error-backpropagation learning. Moreover, rapid modification of the hypothesis is a key point for any unsupervised learning under real-world conditions. By comparison, the input description and the emerging hypothesis should be represented separately at any elementary processing node in the system. Then, every node has to decide locally which state, out of all possible states, it should activate, and which of its representations has to be updated in the process of learning. This has to be done in accordance with its comparison of its state relative to the state of the global system, reflected by the feedforward and feedback inputs. The elementary processing node of networks, which are capable of this kind of hypothetical reasoning, must have a sufficiently complex organization both for local control and knowledge representation in order to act as the required “local agent.” Moreover, this basic organization should be roughly uniform across the system, regardless of its position in the hierarchy.

We propose the *neocortical columnar module* to serve as the required complex elementary processing node. Our system architecture consists of a hierarchy of homogeneous nets whose elementary processing nodes are modular units (Figure 5).

Functionally, the system is composed of two parts that are locally connected at any modular unit: the hypotheses are generated by afferent (sensory) signal flow in a forward processing hierarchy composed of lower subsystems of each modular unit. These hypotheses are firmly established, verified, and updated in a recurrently connected system composed of upper subsystems of each modular unit (Figure 5; for a detailed description of the modular unit architecture see Körner et al. [9]) whose purpose is to recreate the hierarchy of dynamically linked modular units that are expected to be active if the emerging hypothesis is valid. The verification of the hypothesis is performed at any modular unit of the network system by the feedback prediction. Those nodes (modular units) of the



5. Simplified scheme of a neocortical-type artificial neural system architecture. Simplified scheme of bi-directional processing according to our proposed general neocortical-type architecture for self-referential control of recognition and learning. The subsystems at any hierarchic level are composed of modular units that have the same general local control for intra- and inter-modular communication. See Körner et al. [9], [11] for details of architecture and functional implementation.

artificial neural system, which have a forward activated hypothesis and which are consistent with the predictive feedback, can refine this hypothesis, and, at the same time, prime those afferent signal filters, from which the hypothesis originated to stabilize the recognition process. The local hypothesis at one hierarchic level supports those local hypotheses at the lower processing levels that are its components and as such are also part of the consistent global hypothesis (Figure 5). The top-down feedback serves as the “context” for evaluating the subsequent afferent input. The simulation of object recognition in this architecture [11] demonstrated the generation of an initial hypothesis and its subsequent refinement under increasing top-down control of the emerging global hypothesis.

Why the Cortex Needs to Have a Modular, Multilayered Organization

A system like the cortex is in continuous (perpetual) interactions with its environment and responds to any input by a prediction based on its already-acquired knowledge. In neural systems, knowledge is stored in a locally distributed manner, and it modulates the way a sensory input is responded to accordingly within its local vicinity of the global system’s organization. Hence, the comparison of bottom-up processed sensory reality with top-down prediction also has to be performed locally, at each elementary computational node of the system. Moreover, the elementary computational nodes of such a self-referential neural system need to have a firmware architecture, which enables the node to control its behavior within the system’s organization. Depending on its competence to solve a problem, which is measured by its activation resulting from the bottom-up filtering or top-down prediction, this node has to take control over the activation of other nodes in its surrounding, or it has to grant control on its own activation to other nodes in its surrounding. Furthermore, based on the comparison between sensory filtering and top-down prediction, it has to define when it has to learn, and where and what has to be added to the existing representation at this local site. Thus, a self-referential neural system requires complex processing nodes whose architecture reflect the necessary control algorithms for managing these many tasks properly. Since each of the nodes has to cope with the same classes of control

problems, we expect that this architecture of the elementary computational node should constitute a kind of firmware-level organization of self-referential systems.

Flattening the cortex, we find an almost homogeneous layer composed of complex nodes—columnar units—of which the basic internal organization is independent of the information represented there. We propose that this uniform architecture of columnar units does not represent the structure of the knowledge but rather the control structure that forces the system to make those representations.

The firmware-like feature is responsible for the flexibility of processing, especially in the advanced learning capability of an animal with neocortex, and is absent in animals without it. Animals with neocortex dominate the world because they know more about the environment than animals without it, and it is the neocortex that is responsible for this capability [5]. The neocortex should be regarded as the system that makes use of the acquired knowledge to “decompose” the stream of sensory data into meaningful pieces of information. As a result, the expected should be removed from the input description, thus isolating those aspects of the input that are not yet accounted for by the generated internal description [32], [33]. This would explain their enormous selective advantage for speedy and reliable recognition and learning.

In this article, we propose the architecture that can implement the required quality of processing that is unique to living organisms. We attempted to define a general architecture by deriving crucial characteristics of neocortical processing and of its neural organization, which, as we claim, is different from a low-level type of signal filtering and approximation implemented well by formal neural network architectures, composed of simple formal neurons as elementary processing nodes. Neural networks are graphical notations of the classes of algorithms which they support. Hence, different description levels for the encoding of information flow and its inherent control must also be reflected by a proper neural network architecture. Signal filtering and approximation, as we expect to be implemented by the *subcortical* sensory processing stages, are certainly different in their algorithmic structures, compared to the kind of internal simulation at different abstraction levels as we propose for the *cortical* processing stages. The functional

architectures of subcortical and cortical neural processing systems should reflect that difference.

Conclusions

Progress in understanding the brain function under constant interactions with the sensory environment is hampered by inadequate models and theories. Obviously, current models and theories of brain computing still appear irrelevant when they are confronted with real-world problems. We argue that architecture in the brain does not reflect the result of thinking, the ready-made algorithm for solving a problem. Rather it should reflect the control that generates the constraints to select a proper algorithm for a specific problem that is posed by the input—or to create a new one if the application of the previously acquired ones does not provide a sufficient solution. We propose that a value system (based on a genetically imprinted a priori knowledge on coarse behavioral evaluation of sensory input) and neocortical columnar architecture are crucial elements of future artificial neural systems that are expected to emulate the performance of the brain. This should be the case especially for those cognitive tasks that appear easy for animals in their everyday life but turn out to be hopelessly tricky for the current generation of computers.

In order to advance beyond the well-known paradigms of current computational theory, we need a more functional understanding of brain-type computation. As we suggested, the controls that enabled the self-organization of the system may still be the same as those that govern the adaptability and flexibility that we admire when we observe the performance of living beings. Until now, self-referential control algorithms have not been at the focus of attention in neurocomputing and cognitive neuroscience. Models have to replicate more closely the biological original in order to elucidate these intrinsic global control architectures instead of interpreting the performance of neural systems in terms of current computational theory.

Simple architectures generate simple behaviors. Without proper “firmware organization,” artificial neural systems may reasonably perform a certain task (since the respective algorithm is represented by the designed architecture), but they will never achieve the robustness and flexibility that we cherish in the “biological original.” Therefore, morphology and physiology really matter because they

represent the architecture of control we need to understand.

The authors would like to note that this article was first submitted to the magazine in March 1997.



Edgar Körner studied electrical engineering, control engineering, and biomedical cybernetics at the Ilmenau Institute of Technology, Germany. From 1976 to 1984 he served as an assistant professor/senior staff researcher and established the bionics research laboratory at the same university. The research activities included experimental work in neurophysiology and neural systems modeling as well as applied psychophysics and medical expert systems. He received his Dr.-Ing. in biomedical cybernetics in 1977 and the Dr. Sci. in biocybernetics in 1984, both from Ilmenau Institute of Technology. From 1984 to 1987, he joined the Bioholomics Project of JRDC (Tokyo) as a research fellow dealing with brain-like vision systems. Back at the Ilmenau Institute of Technology, he continued research in biological vision and neurofuzzy control systems as an associate professor. In 1988, Dr. Körner was appointed full professor for biocybernetics and head of the Department of Neurocomputing and Cognitive systems. In 1990, he was additionally appointed to head the Department of Neurocomputing at the Institute of Neurobiology and Brain Research at Magdeburg, Germany. In 1992 he moved to Japan to join Honda R&D's Wako Research Center near Tokyo, focusing as a chief scientist on the brain-like computation research. In 1997 he started research in computational neuroscience, evolutionary technology, and cognitive robotics at HONDA R&D Europe, where he serves currently as an executive vice president and head of the Future Technology Research Division.



Gen Matsumoto received his B.S., M.S., and Ph.D. from the Dept. of Physics at the University of Tokyo in 1964, 1966, and 1969, respectively. Dr. Matsumoto was employed at Department of Physics, the University of Tokyo in 1969 and moved in 1971 to the Electrotechnical Laboratory,

MITI (Ministry of International Trade and Industry), working there until 1997 on brain science research in order to develop a brain computer. At first he worked on research of biophysical properties of squid giant neurons at the molecular and cellular levels. He first succeeded in cloning and sequencing squid sodium channel proteins and he first succeeded in maintenance of squid in an aquarium tank in 1975. His interest then focused on the problem of how the brain acquires its algorithm in a self-organizing fashion at a subcellular and cellular level. Dr. Matsumoto was the leader of the MITI Bio-Device Research and Development Project for 1985-1994. For his research objective, he developed an advanced optical imaging system, which was applied to experiments mainly on the rat hippocampus and its surroundings. He developed a neural chip containing 1,000 neurons with 10^6 connections, which can acquire algorithms automatically according to novel learning rules implemented in the chip. In 1994, he was awarded the MITI prize. Dr. Matsumoto was president of the Neural Network Society of Japan (1991-1995) and president of the Biophysical Society of Japan (1995-1997). Since 1997, he has served as the director of the Brainway Group in the Brain Science Institute, RIKEN, to extend his research to a new challenge for revolutionary creation of a brain computer.

Address for Correspondence: Edgar Körner, HONDA R&D EUROPE (Deutschland) GmbH, Future Technology Research Division, Carl-Legien-Str. 30, Offenbach, 63073, Germany. Tel.: +49 69 89011730. Fax: +49 69 89011749. E-mail: Edgar_Koerner@de.hrdev.com.

References

[1] S. Haykin, *Neural Networks*. New York: Macmillan College, 1994.
 [2] G. Hinton, Ed., *Connectionist Symbol Processing*. Cambridge, MA: MIT Press, 1991.
 [3] L. Shastri and V. Ajjanagadde, "From simple associations to systematic reasoning: a connectionist representation of rules, variables and dynamic binding using temporal synchrony," *Behav. Brain Sci.*, vol. 16, pp. 417-494, 1993.
 [4] J.B. Pollack, "On wings of knowledge: a review of Allan Newell's Unified Theories of cognition," *AI*, vol. 59, pp. 355-369, 1993.
 [5] H.B. Barlow, "The coding of sensory messages," in *Current Problems in Animal Behavior*, W.H. Thompson and O.L. Zangwill, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1961, pp. 330-360.
 [6] J.L. Gallant, J. Braun, and D.C. Van Essen, "Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual cortex," *Science*, vol. 259, pp. 100-103, 1993.

[7] A. Arieli, A. Sterkin, A. Grinvald, and A. Aertsen, "Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses," *Science*, vol. 237, pp. 1868-1871, 1996.
 [8] S.P. Vecera and M.J. Farah, "Is visual image segmentation a bottom-up or an interactive process?" *Perception Psychophys.*, vol. 59, pp. 1280-1296, 1997.
 [9] E. Körner, M.O. Gewaltig, U. Körner, A. Richter, and T. Rodemann, "A model of computation in neocortical architecture," *Neural Networks*, vol. 12, no. 7-8, pp. 989-1005, 1999.
 [10] E. Körner and U. Körner, "Robust segmentation of behaviorally relevant events in sensory recognition by synergy of thalamo-cortical and cortico-cortical control," in *Proc. Third Int. Symp. Artif. Life Robot.*, Oita, Japan, Jan. 19-21, 1998, pp. 214-219.
 [11] E. Körner, H. Tsujino, and T. Masutani, "A cortical-type modular neural network for hypothetical reasoning," *Neural Networks*, vol. 10, no. 5, pp. 791-814, 1997.
 [12] D. Marr and A. Vision, *A Computational Investigation into the Human Representation and Processing of Information*. San Francisco, CA: W.H. Freeman, 1982.
 [13] R. Brooks, "A robust layered control system for a mobile robot," *IEEE J. Robot. Automat.*, vol. RA-2, pp. 14-23, 1988.
 [14] E. Körner, U. Körner, and G. Matsumoto, "Top-down selforganization of semantic constraints for knowledge representation in autonomous systems: The role of a value system for autonomous systems," *Bull. Electrotech. Lab.*, vol. 60, no.7, pp. 1-5, 1996.
 [15] D.N. Pandya and E.H. Yeterian, "Prefrontal cortex in relation to other cortical areas in rhesus monkey: Architecture and connections," in *Progress in Brain Research*, vol. 85, H.B.M. Uylings, C.G. Van Eden, J.P.C. De Bruin, M.A. Corner, and M.P.G. Feenstra, Eds. Amsterdam: Elsevier Science, 1990, pp. 63-94.
 [16] B.J. Spiegler and M. Mishkin, "Evidence for the sequential participation of inferior temporal cortex and amygdala in the acquisition of stimulus-reward associations," *Behav. Brain Res.*, vol. 3, pp. 303-317, 1981.
 [17] J. Allman and L. Brothers, "Faces, fear and amygdala," *Nature*, vol. 372, pp. 613-614, 1994.
 [18] R. Adolphs, D. Tranel, H. Damasio, and A. Damasio, "Impaired recognition of emotion in facial expressions following damage to the human amygdala," *Nature*, vol. 372, pp. 669-672, 1994.
 [19] L.M. Romanski and J.E. LeDoux, "Equipotentiality of thalamo-amygdala and thalamo-cortico-amygdala circuits in auditory fear conditioning," *J. Neurosci.*, vol. 12, no. 11, pp. 4501-4059, 1992.
 [20] H.R. Rodman, "Development of inferotemporal cortex in the monkey," *Cerebral Cortex*, vol. 5, pp. 484-498, 1994.
 [21] D.G. Amaral, J.L. Price, A. Pitkaenen, and S.T. Carmichael, "Anatomical organization of the primate amygdaloid complex," in *The Amygdala: Neurobiological Aspects of Emotion, Memory, and Mental Dysfunction*, J.P. Aggleton, Ed. New York: Wiley-Liss, 1992, pp. 1-66.

[22] E. Iwai and M. Yukie, "Amygdalofugal and amygdalopetal connections with modality-specific visual cortical areas in macaques (*Macaca fuscata*, *M. mulatta*, and *M. fascicularis*)," *J. Comparative Neurol.*, vol. 261, pp. 362-387, 1987.

[23] J. Brankack and G. Buszaki, "Hippocampal responses evoked by tooth pulp and acoustic stimulation: Depth profiles and effect of behavior," *Brain Res.*, vol. 378, pp. 303-314, 1986.

[24] O.S. Vinogradova, E.S. Brazhnik, V.F. Kitchigina, and V.S. Stafekhina, "Acetylcholine, theta-rhythm and activity of hippocampal neurons in the rabbit—IV, Sensory stimulation," *Neuroscience*, vol. 53, pp. 993-1007, 1993.

[25] A.A. Ioannides, M.J. Liu, L.C. Liu, P.D. Bamidis, E. Hellstrand, and K.M. Stephan, "Magnetic field tomography of cortical and deep processes: Examples of 'real-time' mapping of averaged and single trial MEG signals," *Int. J. Psychophysiol.*, vol. 20, pp. 161-175, 1995.

[26] J.E. LeDoux, "Emotional memory systems in the brain," *Behav. Brain Res.*, vol. 58, pp. 69-79, 1993.

[27] G.E. Steele and R.E. Weller, "Subcortical connections of subdivisions of inferior temporal cortex in squirrel monkeys," *Visual Neurosci.*, vol. 10, no. 3, pp. 563-583, 1993.

[28] N.M. Weinberger, "Learning-induced changes of auditory receptive fields," *Curr. Opin. Neurobiol.*, vol. 3, pp. 570-577, 1993.

[29] E. Rosch, C.B. Mervis, W.D. Gray, D.M. Johnson, and P. Boyes-Bream, "Basic objects in natural categories," *Cognit. Psychol.*, vol. 8, pp. 382-439, 1976.

[30] P. Cavanagh, "What's up in top-down processing?" in *Representations of Vision*, A. Gorea, Ed. Cambridge, U.K.: Cambridge University Press, 1991, pp. 295-304.

[31] T.W. Picton and D.T. Stuss, "Neurobiology of conscious experience," *Curr. Opin. Neurobiol.*, vol. 4, pp. 256-265, 1994.

[32] H. Barlow, "What is the computational goal of the neocortex?" in *Large-Scale Neuronal Theories of the Brain*, C. Koch and J.L. Davis, Eds. London: MIT Press, 1994, pp. 1-22.

[33] E. Körner, "Autonomous recognition and selforganization of knowledge representation in neural networks. Part 1: From structuring data to

constraint generation by self-referential control," *Holonics*, vol. 4, pp. 3-34, 1994.

[34] S. Ullman, *High-Level Vision*, A Bradford Book. Cambridge, MA: MIT Press, 1996.

[35] S. Celebrini, S. Thorpe, Y. Trotter, and M. Imbert, "Dynamics of orientation coding in area V1 of the awake primate," *Visual Neurosci.*, vol. 10, pp. 811-825, 1993.

[36] M.W. Oram and D.I. Perrett, "Time course of neural responses discriminating different views of the face and the head," *J. Neurophysiol.*, vol. 68, pp. 70-84, 1992.

[37] S. Thorpe, D. Fize, and C. Marlot, "Speed of processing in the human visual system," *Nature*, vol. 381, pp. 520-522, 1996.

[38] H.R. Dinse, "A time-based approach towards cortical functions: neural mechanisms underlying dynamic aspects of information processing before and after postontogenetic plastic processes," *Physica D*, vol. 75, pp. 129-150, 1994.

[39] R.D. Freeman, "Space and time in the central visual pathway," in *Proc. 5th Int. Symp. Bioelectron. Molec. Elect. Dev.*, Okinawa, Japan, Nov. 28-30, 1995, B1-01, pp.1-4.

News Notes



IEEE-USA Focuses on Technology Policy Issues in 2002

IEEE-USA, which represents 235,000 U.S. IEEE members, has refocused its efforts on the most urgent technology policy issues facing the United States in today's turbulent world and is in the process of marshalling its volunteer membership to bring its latest position statements to Congress and federal regulatory agencies.

Earlier this year, the IEEE-USA Board of Directors challenged the IEEE-USA Technology Policy Committee (TPC) to develop position statements advancing the interests of its members, technology, and our country. These positions are to be delivered in person by IEEE-USA volunteers to key congressional committee chairs and members during the summer recess, and continuing in Washington during the rest of 2002. Where appropriate, the issues will also be addressed to regulatory agencies such as the Federal Communications Commission, the Federal Energy Regulatory Commission (FERC), the Food and Drug Administration, and the new Homeland Security Department.

The top priority issues for 2002 include:

- defining a 10-year plan for accelerated broadband deployment, and the consideration of such technologies as ultra-wide band wireless and gigabit ethernet over fiber
- supporting effective measures for security and critical infrastructure protection
- ensuring funding and programs for public health service IT infrastructure, found severely wanting in the aftermath of September 11
- supporting balanced and effective programming of federal investments in R&D, particularly those supporting com-

puter, IT, communications, and networking breakthroughs and advanced development

- ensuring the latest system engineering and technology are brought to bear to improve national aviation safety.

These positions will be emphasized with our elected representatives and executive branch agencies in the remainder of the year.

Other hot issues that will be addressed selectively with appropriate government officials relate to copyright technology controls, the application of trade regulations (ITAR) to technology publications, computer crime, and congressional technology assessment capabilities. For more information, see IEEE-USA's Policy Forum on-line at <http://www.ieeeusa.org/forum>.

The TPC reports to the IEEE-USA Board of Directors through me in my role as IEEE-USA vice president for technology policy and is made up of five volunteer committees covering U.S. policy in communication and information, transportation, energy, medical technology, and research and development.

We are continually soliciting volunteers, whom we will train and ask to join task forces meeting with congressional, executive, and state leaders. We are also soliciting help locally, such as visits with members of Congress in their district offices, placement of TPC-related news in your section newsletter, and/or inclusion of policy discussions in IEEE technical meetings. You can volunteer as part of our Congressional Advocacy Recruitment Effort (CARE) at <http://www.ieeeusa.org/forum/care>, or by contacting IEEE-USA Government Relations Director Chris Brantley at c.brantley@ieee.org.

—Ralph W. Wyndrum Jr.
IEEE-USA Vice President, Technology Policy