

# An egocentric approach to machine intelligence

R. Pienaar

J.J. Kruger

rudolph@enterprise.mikom.csir.co.za  
*Council for Scient. and Indus. Research*  
*PO Box 395*  
*Pretoria, South Africa, 0001*

johannk@tst.co.za  
*Dept. of Electrical and Electronic Eng.*  
*University of Pretoria*  
*Pretoria, South Africa*

## Abstract

*Intimate to the functioning and behavior of intelligent systems is the manner in which information is represented internally. The conventional approach to intelligent system design assumes a particular bias in the manner by which this information is represented. Typically, this is characterized by an “abstract” or “objective” design methodology which holds that intelligence is not a function of the physical nature of the system. Such an approach suffers from several shortcomings, most notably problems relating to scaling and complexity. Recent physiological research, however, has demonstrated that physical bodily form is a fundamental building block in the organization of mammalian cortical structures. Consequently, this article explores such a biologically motivated “subjective” or “egocentric” approach to system design, and demonstrates its utility in a simple robot arm control problem.*

## 1 Introduction

Fundamental to intelligence, be it machine or biological, is the method by which information is organized internally to create internal representations. These internal representations are in turn processed by the system as it satisfies its goal conditions.

The nature of these representations are fundamental to any processing system, comprising its most basic component. This article argues that most conventional approaches to intelligent system design assume a particular “objective” bias in constructing these internal representations, and are consequently plagued with scaling and complexity issues when moving from trivial to “real world” domains. The article will offer a novel design method with a different underlying “subjective” bias, presented as a possible means with which to address the potential shortcomings endemic in the conventional approach.

Of course, the field of robotics spans in and of itself a wide range of views and approaches, from the monolithic philosophy of good old fashioned artificial intelligence, GOFAI, through to the distributed architectures of Brooks’s behavior based robotics [2], [3]. Such behavior based systems are characterized by the interaction of competing sub-systems, each responsible for a particular

motor behavior. Brooks, in addressing issues surrounding scaling such robots to “humanoid”-level systems, defines cognitive robotics, which expand upon behavior based systems in several areas, viz bodily form, motivation, coherence, self adaption, development, and physiological inspiration [4].

Other approaches to scaling robotic intelligence argue for a more ontogenetic perspective [6]. Such views derive from cognitive robotics, developmental psychology, and cognitive neuroscience – arguing for a developmental instead of evolutionary approach. These stem from the assertion that biological intelligences are not born “fully-formed” but develop from immaturity and are shaped by experience, trial and error, goal-orientation, etc. through to adulthood. Development is not evolutionary in the sense of gradual, minor changes to a basically static structure; rather the very framework of this structure is continually refined, broken down, and rebuilt.

An obvious source for cognitive modeling can be found in the physiology of the mammalian cortex itself. Recent work by Kaas [15], [10], [11], [13] has convincingly argued for a re-evaluation of numerous cortical concepts, most notably issues surrounding the multiplicity and extent of sensorimotor mappings, the parallel organization of cortical information pathways, and functional considerations apparent in the nature of internal mappings.

Inspired by basic physiology is work by Bizzi [17], [1], [16] on the Equilibrium Point Hypothesis. This presents a model for muscle and posture control based on the length-tension properties of muscles, and argues for a control-regime quite different to that of conventional robotic approaches. Nonetheless, it lacks an underlying mathematical or signal-flow grounding, and consequently its filtration into conventional robotic control design has been low.

As far as control system design is concerned, Grossberg has made several important contributions, notably the Adaptive Vector Integration To Endpoint (AVITE) model, and offered a framework for learning via motor babbling [7], [8], [9]. Although this AVITE model has been presented in a serial context, its applicability in parallel-processing architectures has not been fully explored (as implied by

physiology) nor have aspects of its information flow been analyzed (as implied by the Equilibrium Point Hypothesis).

## 2 Egocentricity defined

### 2.1 Preliminaries

Egocentricity as a concept addresses the basic nature by which information is organized within a processing system. Logically thus, we need to present some system preliminaries before defining egocentricity.

Consider the interaction mechanism of a processing system with some external environment. Entities within this environment, as well as the environment itself need to be modeled by the system as one or several internal reflections. These internal reflections are in turn processed by the system, in so doing modifying its internal state and/or aspects of the external environment. Any external changes are reflected in a corresponding update of internal reflections; in this manner the interaction process forms a closed loop between the external environment, internal reflections, and back again to the external environment.

The kernel concept lies in the apparently obvious realization that a processing system creates internal reflections. These internal reflections are then manipulated by the system as it interacts with its external environment.

### 2.2 What is egocentricity

Simply stated, egocentricity holds that internal reflections must be constructed from building blocks based on the physical “bodily” form of the system itself. These “body-based” building blocks are in turn manipulated by the system to construct reflections of its external environment.

### 2.3 Why egocentricity

As briefly alluded to earlier, the complexity of “real-world” systems and their method of internal representation pose considerable “scaling” problems. A single jointed arm has a trivial inverse kinematics/dynamics solution; the nature of the inverse solution to a two jointed problem, while more complex, is not necessarily intractable. Adding another joint, not only complicates the problem, but the nature of the solution is different to that of a two jointed limb, and is in turn different to that of a four-jointed system. In short, the mechanism of problem specification and resolution in conventional systems is such that a change to the “real-world” environment often implies a complete re-evaluation and re-structuring of the solution.

Given that processing systems need to construct internal reflections of an external reality, one could argue that the conventional methodology holds such reflections to be completely abstract or “objective”. By this is meant that these internal reflections hold little information about the system that constructed them.

Egocentricity challenges that basic “objective” assertion, offering in its place representations that are “sub-

jective” in the sense that they reflect the physical system-specific attributes.

Fundamentally, egocentricity attempts, by virtue of its derivation from biological concerns, to emulate biological success. Contemporary robots, despite a superficial insect-like appearance, are comparatively clumsy, *inanimate* constructs. Any insights derived from biology can only be to the ultimate benefit of robotic design.

## 3 Physiological support

This section will present an overview analysis of physiological concepts relevant to mammalian cortical organization, demonstrating how egocentricity is a logical integration of this knowledge. First, some basic research in direct support of egocentricity will be presented. Subsequently, two physiologically based models — the Equilibrium Point Hypothesis and VITE structures — will be briefly mentioned. While the basic research of Section 3.1 supports “proof of concept”, the models of Sections 3.2 and 3.3 provide the “means of implementation” that will form the basis of Section 4.

### 3.1 Aspects of mammalian cortical organization

The basic representational building blocks on the cortical surface are *mappings*. Simply stated, a mapping is a continuous area of the cortical surface that represents a body surface in either a motor or sensory context.

Contemporary findings on mammalian physiology can be summarized as:

- Studies on several classes of mammals (primates, carnivores, and some herbivores) have demonstrated the existence of multiple sensory mappings within the visual, auditory, and somatosensory cortices. As a matter of fact, the organization of cortical structures across mammalian species is remarkably similar, implying that fundamental information processing principles are conserved with few variations [11], [13], [14].
- Different body components are reflected *multiply* across the cortex within modalities [15]. There is thus proportionately little non-topographically organized cortex, so much so that the existence of an abstract “association cortex” that is responsible for highest level perception is unlikely.
- From an evolutionary viewpoint, the main difference between more advanced mammals and their more primitive cousins lies in the *number* of sensory/motor mappings, with the number of mappings increasing with more advanced species.
- The *nature* of cortical mappings argues for an organization that reflects functional considerations of the animal’s interaction with its environment [15]. These maps have also been shown to be plastic, adapting to the changing circumstances of a physical bodily component [12].

- The information flow pattern in mammalian cortical structures is not serial, but distinctly parallel. Information propagates from peripheral receptors through the thalamus to the cortex in multiple-component parallel paths, feeding information in different mixes to different cortical representational fields [15] [13].

### 3.2 Equilibrium point hypothesis

The Equilibrium Point Hypothesis is a conceptual model for motor control within biological systems. Given the basic support discussed above, the Equilibrium Point Hypothesis provides insight into the type of initial/final control relationships that can exist between these self-based internal reflections.

#### 3.2.1 Definition

Muscles are arranged about joints in an agonist-antagonist configuration. Consequently, a limb's posture is maintained when the torques exerted by the opposing muscle groups are equal and opposite. The core of the Equilibrium Point Hypothesis rests in the realization that in order to move a limb, it is sufficient for the nervous system to merely specify the pattern of agonist-antagonist equilibrium conditions that would exist across the muscles of the limb *as if the limb were already in the correct position* [5].

#### 3.2.2 Experimental Concerns

In a set of experiments on spinalized frogs (i.e. the surgical disconnection from the brain stem of the spinal cord), Bizzi *et al* has demonstrated convincing support of the Equilibrium Point Hypothesis [17], [1]. Electromyographic analysis indicates that the same general pattern of innervation is supplied to muscles irrespective of the initial orientation, and that it is largely a function of final state.

#### 3.2.3 Egocentric implications

At its root, the Equilibrium Point Hypothesis relies upon the elastic corrective forces of muscles to construct a domain in which movements are processed. These properties are basic and integral to muscles, and also reflect the interaction of muscle systems with their environment. In other words, the Equilibrium Point Hypothesis controls muscle systems based upon a subjective interpretation of their properties and interaction with an external system. Subjective in this context implies a representation derived from operational characteristics of a specific set of muscles unique to a particular animal.

### 3.3 Egocentric control

While the Equilibrium Point Hypothesis sketches the nature of a relationship between internal reflections, Grossberg's VITE/AVITE models can be used as a means to *implement* this relationship. These VITE/AVITE structures are essentially control systems that use the basic structure for both operation and unsupervised learning [7] [8] [9]. The Adaptive Vector Integration To Endpoint (AVITE) is

based on a first order control structure known as the Vector Integration To Endpoint (VITE). The VITE model consists of four basic modules: the TPC, the PPC, a DV unit, and a gated GO signal (see Figure 1).

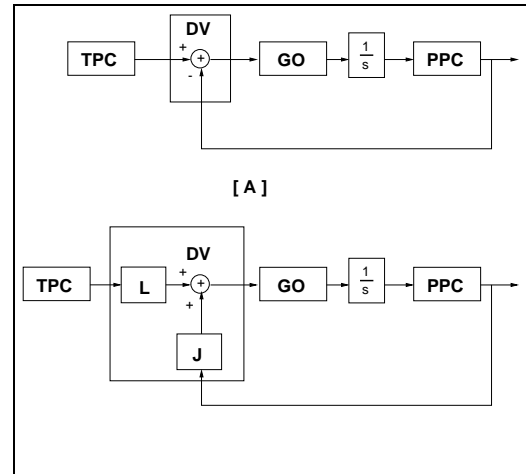


Figure 1: Underlying PI control structure inherent to VITE and AVITE models. [A] PI control loop. [B] DV module expanded to show input filtering matrices.

Essentially, the VITE implements a simple proportional integral feedback control structure to drive the system to a desired state, as shown in Figure 1. Shown within the DV are input filtering  $L$  and feedback filtering  $J$  matrices.

#### 3.3.1 Adaptive learning

The entire purpose of the AVITE model is to present a learning strategy designed to change the weights of the  $L$  matrix in a way that drives the DV to zero. After learning is complete, the DV can only equal zero if the TPC represent the same position or egocentric representation. In effect, therefore, the AVITE generates dimensionally consistent signals from the TPC  $\rightarrow$  DV and from the PPC  $\rightarrow$  DV.

#### 3.3.2 Physiological support

Grossberg [7] [8] has summarized experimental evidence suggesting that the TPC is computed in the parietal cortex, the DV in the motor cortex, and GO signal in the globus pallidus, with the PPC computed in the precentral motor cortex.

#### 3.3.3 Egocentric implications

The VITE models afford us the means by which to associate egocentrically defined state variables with each other by providing a mechanism for associating a particular present representation with a desired target representation. Each internal representation, therefore, has a VITE control structure associated with it.

## 4 Control system implementation

### 4.1 Physical system description

The physical system can be described in two domains: a “real world” external domain, and the internal representation thereof.

As far as the “real world” domain is concerned, the simplest way to describe it is to draw it. Figure 2 shows the major system components. On the top left is a shoulder joint and a limb that articulates with this joint. The shoulder (and hence the attached limb) can rotate around this fulcrum. This limb (referred to as the **arm**) can also extend and contract along its length similar to a radio’s antenna. In this way, the arm is quite unlike most other biological arms. At the other end of the arm is another joint that communicates with the fixed length **hand**. The nature of this joint is such that the hand can rotate 360 degrees around the end of the arm.




Figure 2: The physical robot-arm system. Four “motor” components are present: an arm, a hand, an arm/hand joint, and an arm/shoulder joint

By focusing beyond the detail of AVITE topology, we see that it merely provides the mechanism for linking two entities together. Dropping this unnecessary detail we can define a simpler Object Flow topology that merely shows the topological relation between objects without specifying *how* information flows between them. In other words, we generalize any linking as comprising two *nodes* connected by a single *association map*. These can be simply drawn by representing nodes as open circles, association maps by squares, and by connecting all these together using single arcs, as shown in Figure 3.

An important implication of Figure 3 is that by focusing only on relations between entities and not on an explicit causal flow of information between them, the same general topology can represent an association of  $A \rightarrow B$  or  $B \rightarrow A$  (shown by the selective filling in the small circles on a node’s surface). Information flow in such a topology is set by defining a particular node as a causal source and allowing the effect of this causality to propagate through association maps to other nodes.

Returning to the physical system under consideration, note that individual bodily components can be reflected in individual nodes. Thus, the arm, which has a  $\theta$  and  $\rho$  motor dimension can have two internal reflections, each biased to a particular dimension, reflecting its current state. It can additionally have further reflections in the same egocentric domain reflecting *desired* or target states. We can describe the arm’s egocentric reflections in the following example syntax:  ${}_m T_a$ , where the  $m$  subscript indicates a “motor”, the  $T$  indicates a “target” type (“present” states are likewise denoted by  $P$ ), and the  $a$  subscript indicates an “arm” reflection, whilst  $h$  denotes “hand”. An optional suffix could further specify the *class* of reflection ( $\theta$  or  $\rho$ ).

By referring to Figure 2, we see two entities: an arm and a hand. Additionally, the *entire* workspace forms a visual entity. Note also that another visual entity is defined: the region around the hand.

If we define the set  $\mathcal{N}$  to contain the system’s nodes, we have

$$\mathcal{N} = \{ {}_m T_{a(\theta,\rho)}, {}_m T_{h,v}, T_{a,v}, T_{h,m}, P_{a(\theta,\rho)}, P_{h(\theta),v}, P_{a,v}, P_h \} \quad (1)$$

Whilst nodes provide a logical structure in which to describe bodily reflections, association maps provide the mechanism by which nodes are linked together. We can therefore deduce the necessary association maps given the available nodes and a knowledge of the physical structure of the system. The most obvious association is between a motor target and a motor present (in either the hand or arm modality). Thus, a motor target arm can be linked to a motor present hand by an association module

$${}_m T_a \xrightarrow{T_{ma}} \mathcal{AM}_{ma}^P \rightarrow {}_m P_a \quad (2)$$

Equations of the form of Equation 2 denote relations between nodes and association maps. The usage of the




Figure 3: Demonstrating causality in object flow diagrams. [A] Forward causality (from “A” to “B”) [B] Reverse causality (from “B” to “A”).

right arrow ( $\rightarrow$ ) is to indicate the direction of *forward* causality.

Continuing, in the hand modality we have,

$${}_m T_h \rightarrow {}_{mh}^T \mathcal{A}M_{mh}^P \rightarrow {}_m P_h \quad (3)$$

In Equations 2 and 3, the association map is defined according to the syntax where reading from left to right over the subscripts and superscripts defines the direction of forward causality (in accordance to using the  $\rightarrow$ ). The subscripts indicate the *mode* (motor arm, motor hand, etc), and the superscripts the mode's *type* (target or present in this implementation).

The derivation of the visual association maps develops analogously. Furthermore, if we specify that the processing system relates its hand to its arm in a visual manner, we can deduce that a visual target hand reflection requires information from both the visual target arm node as well as the visual present arm node. Considering also that information from the visual target arm node also influences the visual present arm node, we can relate the major visual nodes as

$$\begin{array}{ccccc} {}_{va}^T \mathcal{A}M_{va}^P & \longrightarrow & {}_v P_a & & \\ \uparrow & & \downarrow & & \\ {}_v T_a & \longrightarrow & {}_{va}^{T,P} \mathcal{A}M_{vh}^T & \longrightarrow & {}_v T_h \end{array} \quad (4)$$

where the association  ${}_{va}^T \mathcal{A}M_{va}^P$  denotes the *motor* nodes and their attendant associations.

In conclusion, if we define the set  $\mathcal{A}$  to contain the association maps, we have

$$\mathcal{A} = \{ {}_{mh}^T \mathcal{A}M_{mh}^P, {}_{ma}^T \mathcal{A}M_{ma}^P, {}_{ma}^T \mathcal{A}M_{va}^P, {}_{ma}^T \mathcal{A}M_{ma}^T, {}_{ma}^P \mathcal{A}M_{va}^P, {}_{vh}^T \mathcal{A}M_{mh}^T, {}_{vh}^P \mathcal{A}M_{vh}^T, {}_{va}^{T,P} \mathcal{A}M_{vh}^T \} \quad (5)$$

## 4.2 Control topology

The final “goal” state of the system can specify not only the fact of the arm “reaching” a particular position in its workspace, but also the final orientation of the hand with respect to this target position. Such an orientation can be implemented as a “memory” reflection. Here, we take a memory to be a particular past state of a bodily component. Within such an egocentric framework, a “memory” is merely an additional node with a fixed internal state, communicating with similar nodes via appropriate association maps. Such a motor hand target memory can be added to the system as shown in Figure 4.

We are now finally in a position to consider the flow of information in this egocentric network and the implications of the Equilibrium Point Hypothesis.

Consider that from a conventional control systems viewpoint a target end state is input to some system which




Figure 4: A memory node. Here a motor target hand memory node is illustrated. It is differentiated from a motor target hand node by the Greek  $\tau$  character.

has an internal present state. Irrespective of any internal mechanisms, the ultimate goal of the control system is the synchronization of target and present states. Relative to the target state space, the present “moves” through a time frame to reach the target. However, relative to the present state’s reference frame, it is static and the target “moves” towards it.

Now, assume that from an arbitrary reference point external to both the present and target states it were possible to physically change the causal flow of information between these states. There are two possible instances, each with the same end goal: the synchronization of present and target. In the one case, causality flows from the target and the effect is the motion of the present. In the other case, causality flows from the present and the effect is the motion of the target.

It is with this realization that we now see the link to the Equilibrium Point Hypothesis. In order to implement this hypothesis, we simply identify a desired end goal, set the corresponding node state accordingly, and assign this node as the source of causality.

Consider Figure 5 which shows the final arm/hand/visual system in its entirety (with the memory node added). Initially, when a visual target becomes active in the visual field,  ${}_v T_a$  acts as the causal source of the network, and information starts to propagate as shown. In order to simplify analysis, we specify that the hand sub-system remains inactive until such a time as the arm is within hand range of the final target.

Now, at some stage, the arm crosses a threshold and is close enough to the target for the hand to reach it. Information flows from  ${}_v T_h$  down through the previously inactive sub-network, supplying the memory node  ${}_m T_h$  and the active node  ${}_m P_h$ . Assume that the memory node contains a desired state for the hand to finally attain. This could be interpreted as the hand touching the target with




Figure 5: The system topology with a motor memory target included. When a visual target becomes active, information initially flows as described by the directed arrows.

a desired final orientation.

By the Equilibrium Point Hypothesis, it should be sufficient for the motor target hand node to be specified as if *it were already in this state*. Practically, this can be realized by reassigning the source of causality to be contained within  $mT_h$ . The flow of information changes throughout the network topology feeding back through  $T_v^P \mathcal{AM}_{v_h}^T$  entering  $vP_a$  and ultimately reaching  $vT_a$ . This has the cascaded result of each node now traveling through its state space so as to realign itself with the new source of causality. Indeed the visual arm *present* state forces the visual arm *target* to change, almost as if the system “sees” where the arm target should move in order to accommodate the desired orientation of the hand.


### 4.3 Results and discussion

For the purposes of this simulation, the system’s boundary conditions are such that the arm and hand start near the bottom left of the workspace (the initial system present state) and end at the top right (the target state). Additionally, the final hand orientation is horizontal so that it “points” at the specified target position, just “touching” it.

In this section, we will analyze the specific case where the association linking the motor target hand ( $mT_h$ ) and the *expected* motor target hand ( $m\tau_h$ ) is a standard proportional integral control structure.

Additionally, the flow of information when the expectation node becomes a source is such that the visual *present arm* is driven by the motor *target hand*, as shown in Equation 6. Causality flows ultimately back to the visual target

arm, which moves now to track the visual present arm – a prime example of how cause and effect nodes have been reversed.



The system operates in two distinct phases, corresponding with the two sources of causality. At first the visual target arm is the source of causality; yet once the arm is sufficiently close to the target, the source of causality switches to the expected motor target hand, ( $m\tau_h$ ). Subsequent system behavior retains the source of causality here as shown in Figure 6 demonstrating the development of hand- $\theta$  reflections. Figures 7 and 8 respectively illustrate the  $\theta$ - and  $\rho$ -arm modal trajectories.




Figure 6: Motor hand-based  $\theta$  trajectories. The expected value is  $\theta = 0$ , and when active, attracts the motor target hand trajectory, which in turn attracts the motor present hand trajectory.

Consider the  $mT_h$  trajectory. The first three seconds (approximately) of simulation time correspond to the period during which the hand reflections are inactive or dormant. No information flows in the hand-based branch of the network topology, and the hand itself can be considered in a “resting” state. The  $mP_h$   $\theta$  state is “pointing down”, and  $mT_h$  is in a random dormant state (in this case  $mT_h(\theta_0) = 1$ ).

As the arm approaches the target point, the hand-based visual target ( $vT_h$ ) will at some stage become stimulated, containing an active point on the boundary of its




Figure 7: The motor  $\theta$  state trajectory. The independent coordinate ( $x$ -axis) is time in seconds.

retinal map. This corresponds to the moment when the final target point is within range of the hand.

At this time, the position of the system's final target relative to the center of the hand's visual space can be determined. For the current simulation, this corresponds to the step response of  $mT_h$  from 1 to -6. Simultaneously, the memory node becomes a causal source, and consequently starts driving  $mT_h$  to 0 (there is a brief delay before  $mT_h$  is influenced – this is due to internal conditions specific to the software simulation).

The organization of the hand-based modes are such that  $mT_h$  in turn drives  $mP_h$ . This is reflected in the response of  $mP_h$  as it tracks from its initial position of  $mP_h(\theta_0) = 6$  to its steady state position of  $mP_h(\theta_{ss}) = 0$ .

Simply stated, a related chain of cause and effect can be defined: initially the motor memory becomes active. This is subsequently tracked by the motor target, which in its turn is tracked by the motor present. Since all of these modes communicate via first order control structures, and considering that from the network topology the motor present is two association maps away from the motor target memory, one could postulate that the motor present trajectory would display second order characteristics, which is confirmed in Figure 6.

Related behavior is also evident in the trajectories of the arm-based motor modes. Consider Figure 8 which shows the expected arm-based  $\rho$  plots. As information flows outward from the motor target memory node, an information flow path leads back through the visual target hand ( $vT_h$ ), and from there flows into the visual present arm ( $vP_a$ ) before reaching the present arm motor nodes.

The approximately first three seconds of Figure 8 shows  $mP_\rho$  following first order convergence as it tracks  $mT_\rho$  due to the causal source in  $vT_a$ . A slight "blip" is evident when the rate of convergence has slowed to the point where topological considerations of the underlying  $\rho$ -mask sur-




Figure 8: The motor  $\rho$  state trajectory. The independent coordinate ( $x$ -axis) is time in seconds.

face come into play. Hereafter, the source of causality shifts from  $vT_a$  to  $mT_h$  and consequently  $mP_\rho$  becomes driven by information flowing from  $vT_h$  and  $vP_a$ . The step-like behavior of  $mP_\rho$  is due to the quantization of the  $vT_h$  and  $vP_a$  visual modes.

Furthermore,  $vT_\rho$  now tracks  $vP_\rho$ , converging as shown. Causality in a very real sense flows backward as the *target* now adapts itself to be consistent with the *present*. When the simulation has ended and one considers the end states without knowledge of how the system evolved, one might be tempted to postulate that the visual target arm's entire trajectory was calculated from the beginning and drove the present trajectory throughout. In reality though, the visual target drove the system for a while, and then causality shifted completely such that for the last part, the target was driven by the present. The end result is the same, although the approach followed here is considerably simpler to implement and requires no prior complex calculations.

## 5 Conclusion

This paper addressed issues pertaining to the processing of information within intelligent systems. Specifically, the mechanisms by which internal reflections are constructed, as well as their relation to the processing system were explored.

In effect, current physiological research was used to create a design framework within which a new processing model was created. This model, or egocentric construct, held that internal reflections of environmental entities are strongly based on the system's physical "bodily" form and its subjective interaction with an environment. Simply stated, egocentricity at its root maintains that reflections of a system's physical body form the basis of all processing. This basic precept is strongly supported by relatively recent physiological research on mammalian cortical or-

ganization.

An example simulated robot controller was presented. This example problem demonstrated how relatively simple it is to accommodate increasing external complexity by merely expanding the processing topology accordingly (the addition of a hand to the arm, and the addition of a hand memory to the whole system).

Alluding back to the introduction, we propose that such work can contribute to the cognitive robotics of Brooks as well as the ontogenetic perspective of Kemp by providing a definite link back to more physiologically inspired research, explicitly demonstrating the importance of “bodily” form in the construction of a processing topology. More importantly perhaps, egocentricity provides a common philosophical framework within which to cast the somewhat disparate knowledge “sources” that Brooks has drawn upon to develop cognitive robotics.

In a similar vein, egocentricity allows the contributions of Kaas *et al* to be efficiently integrated into a systems-based philosophy, thus opening the door and allowing basic physiological work to make a more fundamental contribution to systems engineering.

Likewise, the Equilibrium Point Hypothesis can be applied to control systems as was shown, implying the powerful concept of “fluid” causality whereby different nodes within a processing topology can function as causal sources at different times.

The VITE and AVITE control structures proposed by Grossberg were originally presented in a serial context. Here, they were applied in parallel and used to associate *egocentric* data structures — demonstrating their usefulness in a problem domain possibly quite removed from their original home. Combined with the implications of the Equilibrium Point Hypothesis, they formed the basic mathematical building blocks of the entire processing system.

Finally, egocentricity within this introductory paper is presented as a possible philosophy uniting cortical physiology and robotic engineering. It is hoped that future research within such an egocentric context can address problems relating to machine learning, i.e. the self-generation of a processing topology. Moreover, it is hoped that egocentricity can provide a fresh perspective on intelligent systems and provide the means by which new insights into not only system design but also physiological processing can be developed.

## References

- [1] E. Bizzi and F. A. Mussa-Ivaldi. Does the nervous system use equilibrium-point control to guide single and multiple joint movements? *Behavioral and brain sciences*, 15:603–613, 1992.
- [2] R. A. Brooks. Intelligence without reason. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 569–595, 1991.
- [3] R. A. Brooks and L.A. Stein. Building brains for bodies. Technical Report Memo 1439, MIT Artificial Intelligence Laboratory, 1993.
- [4] R.A. Brooks. From earwigs to humans. Technical report, MIT Artificial Intelligence Laboratory, 1997.
- [5] A. G. Feldman. Central and reflex movements of motor control. *Biofizika*, 11:667, 1966.
- [6] C. B. Ferrell and C. C. Kemp. An ontogenetic perspective to scaling sensorimotor intelligence. Technical report, MIT Artificial Intelligence Laboratory, 1995.
- [7] P. Gaudiano and S. Grossberg. Vector associative maps: Unsupervised real-time error-based learning and control of movement trajectories. *Neural Networks*, 4:147–183, 1991.
- [8] P. Gaudiano and S. Grossberg. Adaptive vector integration to endpoint: Self-organizing neural circuits for control of planned movement trajectories. *Human Movement Science*, 11:141–155, 1992.
- [9] S. Grossberg, F. Guenther, D. Bullock, and D. Greve. Neural representations for sensory-motor control, ii: Learning a head-centered visuomotor representation of a 3-d target position. *Neural Networks*, 6:43–67, 1993.
- [10] J. H. Kaas. The organization of the neocortex in mammals: Implications for theories of brain function. *Annual Review of Psychology*, 38:129–151, 1987.
- [11] J. H. Kaas. How the sensory cortex is subdivided in mammals: Implications for the studies of the prefrontal cortex. *Progress in Brain Research*, 85:3–11, 1990.
- [12] J. H. Kaas. Plasticity of sensory and motor maps in adult mammals. *Annual review of neuroscience*, 14:137–167, 1991.
- [13] J. H. Kaas. The functional organization of somatosensory cortex in primates. *Annals of Anatomy-Anatomischer Anzeiger*, 175:509–518, 1993.
- [14] L. A. Krubitzer and J. H. Kaas. Cortical integration of parallel pathways in the visual system of primates. *Brain Research*, 478:161–165, 1989.
- [15] M. M. Merzenich and J. H. Kaas. Principles of organization of sensory-perceptual systems in mammals. *Progress in Psychobiology and Physiological Psychology*, 9:1–42, 1980.
- [16] F. A. Mussa-Ivaldi and E. Bizzi. Linear combinations of primitives in vertebrate motor control. *Proceedings of the National Academy of Science of the USA*, 91:7534–7538, August 1994.
- [17] F. A. Mussa-Ivaldi, S. F. Giszter, and E. Bizzi. Motor-Space Coding in the Central Nervous System. *Cold Spring Harbor Symposia on Quantitative Biology*, 60:827–835, 1990.