Proceedings of the 2007 IEEE Symposium on
Artificial Life (CI-ALife 2007)

# Control-theoretic Synthesis of Hierarchical Dynamics for Embodied Cognition in Autonomous Robots

Nebu John Mathai and Takis Zourntos
Department of Electrical and Computer Engineering
Texas A&M University
College Station, Texas 77843-3128
Email: {mathai,takis}@neo.tamu.edu

***Abstract*— We consider the design of a continuous hierarchical dynamic system architecture to endow cognition in an autonomous mobile robot with bounded resources. The hierarchy is decomposed on the basis of behavior with lower levels being concerned with motion control, and upper levels being concerned with ensuring overall progress. Our architecture consists of a composition of regulators, all of which have access to sensor data and which control problems relating to the agent's motion in the world (i.e., its velocity, position, and path); it hence realizes a physically embodied cognitive system. To synthesize the regulators, we employ a plant-controller co-design methodology; after constructing a plant that models the salient features of the environment relevant to the behavioral level under consideration, we apply rigorous nonlinear control-theoretic toolsets to synthesize a corresponding controller. An interesting side-effect of our synthesis is the emergence of normalized and unnormalized radial basis functions in our control law. Simulation results and animations (available at the URL indicated in Section V) illustrating the life-like behavior of the system are provided. We highlight the emergence of *satisficing* intelligent behavior by our cognitive architecture and show how this satisfies Bedau's definition of *weak emergence*.**

## I. Introduction

Life and cognition are strongly interrelated. Sommerhoff's early attempt at creating the foundations for an analytical biology [1] sought to explain the apparent "purposiveness" of life. The cybernetics movement sought to isolate the control and communications mechanisms underlying artificial and natural autonomous, goal-directed systems—mechanisms that were felt could be applied to "all of the exteriorally directed [i.e., cognitive] activities of an organism" [2]. Maturana and Varela in their work on autopoietic systems [3] indicated the direct connection between life and cognition through the statement "living systems are cognitive systems, and living is a process of cognition." Finally, work on embodied cognition and the "artificial life route to artificial intelligence" [4] make strong cases for their work based on insight gained from considering[1] *living* autonomous, cognitive systems.

N.J.Mathai acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) PGS D Scholarship.

[1]e.g., through ethology or neuroscience

The view that embodiment—i.e., strong coupling between agent and environment—plays an important role in cognition has a long history in the cybernetics, robotics, artificial intelligence (AI), and artificial life communities [2], [3], [5]–[13]. Ziemke [13] surveys the differing views on embodiment, noting that despite the acknowledgment of the vital role of embodiment in cognition, there are a diversity of notions of what embodiment is. In particular, *physical embodiment* is defined as the strong coupling of an agent to the environment via sensors and actuators. We find that this definition is relevant to robotics and reflects early cybernetics-oriented positions, such as that of Albus [9] who takes the position that the "sensory-interactive goal-directed motor system is not simply an appendage to the intellect, but is rather the substrate in which intelligence evolved." Furthermore, he suggests that the cognitive structures used for so-called "low-level" motor behavior may be, under slight modification, applied to endowing higher-level faculties. Matarić looks at this idea in [11].

Taking a view similar to [9] we develop a physically embodied cognitive architecture in which the primitive motif of a regulatory system is replicated at all levels. Dynamical systems theory and cybernetics provide a natural framework to deal with such coupled agent-environment systems [14]. In particular, nonlinear control theory (a modern descendant of cybernetics) yields rigorous synthesis toolsets that can be used to realize dynamic system architectures for goal-directed systems. We are also motivated towards control-theoretic approaches by the promise of an analog[2] system architecture. Lumelsky [15] suggests that continuous-time "dynamic" architectures are a natural substrate for robotics problems in which an agent with bounded resources must attain a global goal with only local information. In addition, the natural parallelism of analog "computation" systems, combined with the promise of economical analog implementation made possible by neuromorphic engineering [16] motivates us to seek system formulations that are amenable to analog realization.

[2]By "analog" systems we mean those that operate on analog signals, which are continuous in amplitude (and either continuous or discrete in time). In contrast, digital signals are discrete in both time and amplitude.

1-4244-0701-X/07/$20.00 ©2007 IEEE

### A. Contributions

In this work we focus on the development of embodied cognition to endow a resource-bounded agent with the ability to autonomously navigate through an obstacle-ridden spatial environment to find a target. We mate control-theoretic toolsets with a hierarchical dynamic system architecture to synthesize these cognitive faculties. An interesting by-product of our approach is the emergence of normalized and unnormalized radial basis functions in the control law.

Our approach to control is novel in that we do not seek the asymptotic stability of the system. Indeed, we find that by relaxing the aggressiveness of the control law to ultimate boundedness [17], *satisficing intelligence* [18]–[20] emerges [21]. We demonstrate this through simulation results, and show how this emergence is in line with Bedau's definition of *weak emergence*.

## II. Development

We consider an environment, represented in Euclidean space, $\mathcal{R}^3$, in which the target, obstacles and initial robot position are located within a compact subset of $\mathcal{R}^3$ denoted by $\mathcal{C} \subset \mathcal{R}^3$. The $N_o$ obstacles, $N_o \in \mathcal{Z}$, $0 \leq N_o < \infty$, denoted by $\Omega_i$, $i = 1, \ldots, N_o$, are closed simply connected sets satisfying $\Omega_i \subset \mathcal{C}$ for all $i$ and $\Omega_i \bigcap \Omega_j = \emptyset$ for all $i \neq j$. The dimensionless, stationary target, denoted by $T$, is located at position $\boldsymbol{p}_T \in \left(\mathcal{C} - \bigcup_{i=1}^{N_o} \Omega_i\right)$. We assume that there is some $\rho > 0$ such that the $\rho$-neighborhood about $\boldsymbol{p}_T$, $B_T := \{\boldsymbol{z} \in \mathcal{R}^3 : \|\boldsymbol{z} - \boldsymbol{p}_T\| \leq \rho\}$, satisfies $B_T \subset \mathcal{C}$ and $B_T \bigcap \Omega_i = \emptyset$, $i = 1, \ldots, N_o$, where $\| \cdot \|$ denotes the Euclidean norm.

The robot, dimensionless but with a specific orientation, is denoted by $M$ and is located at position $\boldsymbol{p}_M : \mathcal{R} \mapsto \mathcal{R}^3$, which for the initial time $t = t_0$ is $\boldsymbol{p}_{M,0} := \boldsymbol{p}_M(t_0)$. The robot has a local frame of reference.

The basic problem we address is the real-time determination of a path from the starting position $\boldsymbol{p}_{M,0}$ to a boundary point of $B_T$, assuming no knowledge of the environment. So that a solution is possible, we assume the existence of a collision-free path, defined as $\boldsymbol{p}_s : \mathcal{R} \to \mathcal{R}^3$, with $\boldsymbol{p}_s(t_0) = \boldsymbol{p}_{M,0}$ and $\boldsymbol{p}_s(t_f) \in \partial B_T$, where $t_f$ satisfies $t_0 \leq t_f < \infty$, such that for any $i$ and for any $t$, $\boldsymbol{p}_s(t) \notin \Omega_i$.

### A. Actuation

We attach a three-dimensional local coordinate system (with orthogonal axes $x_{1L}, x_{2L}, x_{3L}$) to the robot such that the robot is at the origin, and the robot is oriented so that its forward direction is along the positive $x_{1L}$ axis. It has one actuator to effect forward and reverse translational motion in the direction of the $x_{1L}$ axis, and two actuators that enable (independent) rotations about the remaining axes, $x_{2L}$ and $x_{3L}$. These three actuators allow the robot to achieve any orientation and position in $\mathcal{R}^3$.

### B. Sensing

In this section, we provide mathematical models for the sensors; note, however, that these sensor models are representative only, and that any global information used in these models is not necessary for our proposed robot control scheme. For example, local obstacle detection could, in practice, be implemented by means of an ultrasonic device rather than by the computations implied in (1) and (2), below.

The robot perceives the environment through four sensors, with vector-valued sensor outputs denoted as $\boldsymbol{\sigma}_k$, $k = 0, 1, 2$, and a scalar-valued sensor output $\sigma_3$. We define the distance of the robot from the target as $\rho_T(t) := \|\boldsymbol{p}_M(t) - \boldsymbol{p}_T\|$. The influence of obstacles is given in the local spherical coordinates of the robot as a disturbance opposing any forward translation:

$$\boldsymbol{\sigma}_0(t) = \begin{bmatrix} \sigma_{0,1} \\ \sigma_{0,2} \\ \sigma_{0,3} \end{bmatrix} := \begin{bmatrix} -l_O[d(t)] \\ 0 \\ 0 \end{bmatrix} \quad (1)$$

where $l_O : \mathcal{R} \mapsto \mathcal{R}$ is a class-$\mathcal{L}$ function, i.e., a continuous positive strictly decreasing function such that $l_O(0) > 0$ and $\lim_{d \to \infty} l_O(d) = 0$, $d(t) := \min_{\forall \boldsymbol{w} \in \boldsymbol{W}} \|\boldsymbol{p}_M(t) - \boldsymbol{w}\|$,

$$\boldsymbol{W} := \boldsymbol{S}(\boldsymbol{p}_M(t), \theta_{M,2}(t), \theta_{M,3}(t)) \bigcap \left( \bigcup_{i=1}^{N_o} \Omega_i \right) \quad (2)$$

and $\boldsymbol{S}(\boldsymbol{p}_M(t), \theta_{M,2}(t), \theta_{M,3}(t))$ (where $\theta_{M,2}(t)$ and $\theta_{M,3}(t)$ denote the orientation of the robot with respect to the global coordinate system) is the set of all points in some sector originating from $\boldsymbol{p}_M(t)$ that includes the positive $x_{1L}$ axis.

The sensor signal $\sigma_3$ is given by:

$$\sigma_3 = h_0[l_T(\rho_T)], \quad (3)$$

where $h_0 : \mathcal{R} \mapsto \mathcal{R}$ is a hysteresis function of $l_T(\rho_T)$ as shown in Figure 1, and $l_T : \mathcal{R} \mapsto \mathcal{R}$ is a class-$\mathcal{L}$ function. The hysteresis thresholds, $T_L$ and $T_H$, satisfy $T_L < T_H$.

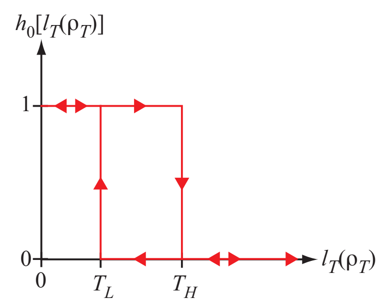


Fig. 1. Hysteresis characteristic for sensor output, $\sigma_3$.

The target sensor output, $\boldsymbol{\sigma}_1$, is given with respect to the local coordinate system of the robot as:

$$\boldsymbol{\sigma}_1(t) = \begin{bmatrix} \sigma_{1,1} \\ \sigma_{1,2} \\ \sigma_{1,3} \end{bmatrix} := \begin{bmatrix} r_T(t) \\ \theta_{T,2}(t) \\ \theta_{T,3}(t) \end{bmatrix} \quad (4)$$

and provides the relative position of the target in spherical coordinates. If $\sigma_3 = 1$, then $\boldsymbol{\sigma}_1$ is considered to be "invalid" and is not used for the purposes of motion control (since, in practical terms, the target is out of sensor range).

Finally, the "inertial guidance" sensor produces the signal $\boldsymbol{\sigma}_2$ composed of the translational speed ($v$) and the rotational speeds ($\omega_2, \omega_3$) of the robot:

$$\boldsymbol{\sigma}_2 = \begin{bmatrix} \sigma_{2,1} \\ \sigma_{2,2} \\ \sigma_{2,3} \end{bmatrix} := \begin{bmatrix} v \\ \omega_2 \\ \omega_3 \end{bmatrix}. \quad (5)$$

## III. Hierarchical Dynamic Architecture

Dynamic hierarchical organization is a fundamental observed characteristic of living organisms whose importance has been addressed in the literature [22]–[25]. In fact, [22] poses the synthesis of dynamical hierarchies at all scales as one of fourteen crucial open problems of artificial life. We provide a brief survey of some representative physically-embodied hierarchical dynamic architectures.

An early formulation of a continuous-time hierarchical dynamic architecture was Ashby's *ultrastable* system [5]. A two level hierarchy was specified consisting of a lower-level "reacting" part strongly coupled to the environment, and a higher-level system operating on a slower time scale that regulated the lower-level system based on environmental feedback. In the practical realization of the ultrastable system—the Homeostat—the higher-level system possessed the ability to *search* for successful controls to regulate the lower-level system.

Albus proposed a cascaded continuous-time architecture [9] where the output of one level becomes the input to the adjacent lower level. Sensory feedback from the environment entered all levels of the hierarchy, with higher levels possibly using abstractions of lower-level senses (e.g., sensory information from which pertinent features have been extracted). A useful architectural insight from Albus is the suggestion that "the same type of anatomical components which are used in the lower and mid levels of the control hierarchy to produce sensory interactive motor behavior may . . . be used at the upper levels of the same hierarchy to plan and solve problems."

Finally, Brooks [10] applies a *vertical* behavior-based decomposition to generate the *subsumption* architecture. Here, each "level of competence" has access to sensors and actuators and forms a complete robot controller on its own.

### A. Our Approach

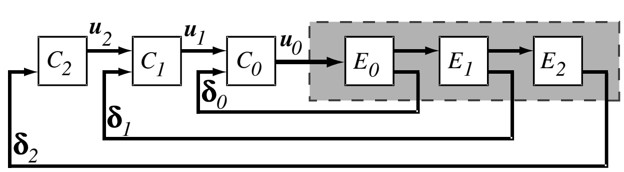


Fig. 2. Hierarchy of cascaded regulators, $C_i$, interacting with the environment, $E = \{E_0, E_1, E_2\}$.

We describe our hierarchical architecture (illustrated in Figure 2) as a cascade of regulating systems where the agent's cognitive component (controller, $C_i$) seeks to regulate its sensory perception of the environment ($E$). We develop the functional complexity of the cognition in a piecemeal fashion, with each level introducing a new fundamental mode of behavior. Our formulation is physically embodied, with each level of the hierarchy receiving information from the environment and actuating change on the basis of this information.

The lowest level, level zero, is concerned with causing motion. The regulation problem here is the selection of acceleration commands for the agent's actuators to achieve a velocity objective. At this level, beyond physical limitations posed by the actuators, there is very little constraint on the types of motion that are allowed.

Level one constrains the agent's motion to ensure it avoids obstacles while tracking a target or a reference velocity objective. We pose the regulation problem as the tracking of an objective (either target-tracking or velocity-tracking) while suppressing obstacles; in Section IV-B we see that this can be naturally formulated through the topology of a closed-loop feedback system.

We note that level one presents the agent with potentially conflicting goals (e.g., the problem where suppressing the obstacle involves moving away from the target), as well as situations where the agent is not motivated to move (e.g., the target is out of range). Both situations can lead to lack of progress by the agent; in the former case, limit cycles are possible as the agent attempts to resolve its conflict, and in the latter case, the agent stays still waiting for stimulus. Level two addresses this by excluding these cases—we introduce searching behavior in which the agent engages in "trial-and-error goal-seeking" [9] to bring itself back into a domain where the level one controller can properly regulate—this is reminiscent of Ashby's ultrastable system. Hence, at this level the type of path executed by the agent is constrained.

Table I summarizes the hierarchical organization of our architecture showing the task decomposition of each level and the scale of operation.

| Level | Behavior | Scale |
|---|---|---|
| 0 | motion | velocity |
| 1 | velocity-tracking | velocity, position |
| | target-tracking (taxis) | position |
| | obstacle avoidance | position |
| 2 | searching | path |

TABLE I
Proposed hierarchical organization.

## IV. Synthesis

We now look at each level of the hierarchy and design a controller for that level using a methodology of plant-controller co-design. First, we construct a system (the plant) that models the salient dynamics of the problem posed by the hierarchical level under consideration. In creating the plant model, we make some simplifying assumptions about the external world; these assumptions are supported by:

- the subsumption architecture of our cognition (i.e., that each level provides a cognitive skill that upper levels can exploit)

- the separation of time scales between each level of the hierarchy (i.e., each level operates at a faster time scale that the one above it, as seen in Table I)

Hence, level zero (operating at the fastest time scale) assumes that the actuators can instantaneously achieve acceleration commands, level one assumes that level zero enables instantaneous tracking of velocity commands, and level two (operating at the slowest time scale) assumes that level one provides obstacle avoidance while executing commanded paths.

With a plant model, we have a quantitative specification of the problem enabling the application of rigorous control-theoretic toolsets (a combination of backstepping [26], [27] and Lyapunov synthesis) to design a solver (i.e., the regulator) for the problem. Our premise is that the solution of the problem in the plant-controller domain will yield a solution for the associated problem in the real-world domain. Further, due to the strong coupling with the environment (due to the sensor and actuator signals that bring information from and to the environment, respectively) the cognition that emerges from the composition of these controllers will be physically embodied.

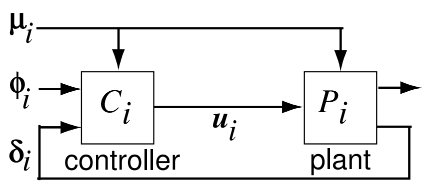


Fig. 3. Basic plant-controller motif.

The basic plant-controller motif that we employ is shown in Figure 3 for the $i$-th level of hierarchy, where $P_i$ denotes the plant model, $C_i$ denotes the associated controller, $\boldsymbol{\eta}_i$ denotes the state of the plant model, $\boldsymbol{\phi}_i$ denotes the reference value to the controller, $\boldsymbol{\delta}_i$ denotes the sensory perception of the environment that is relevant to the cognitive skill for which $C_i$ is being designed, $\boldsymbol{\mu}_i$ denotes parameters to the model, and $\boldsymbol{u}_i$ denotes the controller output (the "control"). We note that $P_i$ models the effect of all controllers downstream from $C_i$ (i.e., $C_0, \ldots, C_{i-1}$) due to the cascaded hierarchy of Figure 2. The general control problem then is to design a control law for $C_i$ that will drive $\boldsymbol{u}_i$ to actuate change that will cause the error $\boldsymbol{e}_i := \boldsymbol{\phi}_i - \boldsymbol{\delta}_i$ to go to zero (i.e., $\lim_{t\to\infty} \boldsymbol{\delta}_i \to \boldsymbol{\phi}_i$).

## A. Level Zero

For the motion causation problem, we assume an actuator that is able to instantaneously respond to acceleration commands:

$$\boldsymbol{u}_0 = \begin{bmatrix} u_{0,1} \\ u_{0,2} \\ u_{0,3} \end{bmatrix} := \begin{bmatrix} a \\ \alpha_2 \\ \alpha_3 \end{bmatrix} \tag{6}$$

(where $a$ is translational acceleration, and $\alpha_2$ and $\alpha_3$ are rotational accelerations about the $x_2$ axis and $x_3$ axis, respectively) Then, with $\boldsymbol{\eta}_0 = \boldsymbol{\sigma}_2$ we construct a model of how the inertial guidance sensor responds to $\boldsymbol{u}_0$:

$$P_0 : \begin{cases} \dot{\boldsymbol{\eta}}_0 = \boldsymbol{u}_0 \\ \boldsymbol{\delta}_0 = \boldsymbol{\eta}_0 \end{cases} \tag{7}$$

Now, we synthesize a controller that will track the reference velocity command:

$$\boldsymbol{\phi}_0 = \begin{bmatrix} \phi_{0,1} \\ \phi_{0,2} \\ \phi_{0,3} \end{bmatrix} := \begin{bmatrix} \phi_{0,v} \\ \phi_{0,\omega_2} \\ \phi_{0,\omega_3} \end{bmatrix} \tag{8}$$

(where $\phi_{0,v}$ is the command translational speed, and $\phi_{0,\omega_2}$, $\phi_{0,\omega_3}$ are the command rotational speeds). Define the tracking error:

$$\boldsymbol{e}_0 = \begin{bmatrix} e_{0,1} \\ e_{0,2} \\ e_{0,3} \end{bmatrix} := \boldsymbol{\phi}_0 - \boldsymbol{\delta}_0 \tag{9}$$

and the scaler-valued function:

$$V_0 := \frac{1}{2}\boldsymbol{e}_0^T \boldsymbol{e}_0 \tag{10}$$

which is positive-definite with respect to $\boldsymbol{e}_0$. Differentiating (10) with respect to time, we obtain:

$$\dot{V}_0 = \boldsymbol{e}_0^T(\dot{\boldsymbol{\phi}}_0 - \boldsymbol{u}_0). \tag{11}$$

Let $K_0$ be a positive-definite matrix. Then setting $\boldsymbol{u}_0 = \dot{\boldsymbol{\phi}}_0 + K_0\boldsymbol{e}_0$ makes (11) negative-definite with respect to $\boldsymbol{e}_0$, and hence $\lim_{t\to\infty} \boldsymbol{\delta}_0 \to \boldsymbol{\phi}_0$ with the rate of convergence set by $K_0$.

## B. Level One

We begin by recasting the problem of obstacle avoidance and sensor tracking into a multivariable feedback control problem. Figure 4 illustrates a refinement to the motif of Figure 3 where we introduce a signal *suppression* channel (located where $\boldsymbol{\sigma}_0$ is injected) in addition to the tracking channel (as before, where $\boldsymbol{\phi}_1$ and $\boldsymbol{\delta}_1$ are injected).

The parameter $\mu_1 \in \{0, 1\}$ selects whether the system should track the target sensor (i.e., engage in taxis) or the inertial guidance sensor. The block $P_1$ models sensor outputs

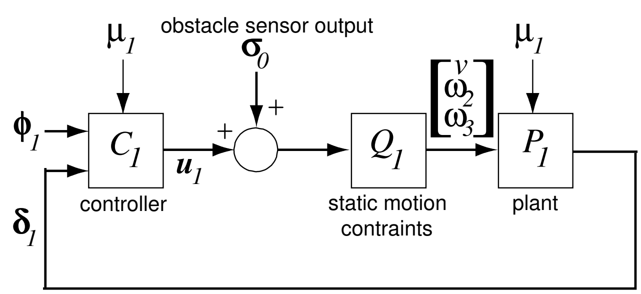


Fig. 4. Control system topology for level one motion control.

$\boldsymbol{\sigma}_1$ and $\boldsymbol{\sigma}_2$ as the vehicle moves through the environment. The dynamic model of $P_1$ can be derived using straightforward geometry and is given by:

$$P_1 : \begin{cases} \dot{\boldsymbol{\eta}}_1 = \boldsymbol{B}_1(\boldsymbol{\eta}_1) \begin{bmatrix} v \\ \omega_2 \\ \omega_3 \end{bmatrix} \\ \boldsymbol{\sigma}_1 = \boldsymbol{\eta}_1 \\ \boldsymbol{\sigma}_2 = \begin{bmatrix} v \\ \omega_2 \\ \omega_3 \end{bmatrix} \\ \boldsymbol{\delta}_1 = (1-\mu_1)\boldsymbol{\sigma}_1 + \mu_1 \boldsymbol{\sigma}_2 \end{cases} \quad (12)$$

where $\boldsymbol{\eta}_1 = \begin{bmatrix} r_T \\ \theta_{2,T} \\ \theta_{3,T} \end{bmatrix}$,

$$\boldsymbol{B}_1(\boldsymbol{\eta}_1) := \begin{bmatrix} -\sin(\theta_{2,T})\cos(\theta_{3,T}) & 0 & 0 \\ -\frac{\cos(\theta_{2,T})\cos(\theta_{3,T})}{r_T} & -\cos(\theta_{3,T}) & 0 \\ \frac{\sin(\theta_{3,T})}{r_T \sin(\theta_{2,T})} & \frac{|\sin(\theta_{3,T})|\cos(\theta_{2,T})}{\sin(\theta_{2,T})} & -1 \end{bmatrix}$$

For simplicity, we assume that the robot instantaneously achieves any input translational or rotational velocity (which can be approximately achieved through a "high-gain" selection of $K_0$ in the level zero controller), but may be subject to external bounded actuator disturbances (either stemming from the environment or system constraints such as quantization of speed). Hence, the block $Q_1$ is free of dynamics, and models static nonlinearities in the actuators (e.g., speed quantization, deadzone characteristics, and other memoryless functions).

The controller block, $C_1$, constitutes the agent's "intelligence." The formulation of Figure 4 indicates that the information about the environment available to the controller is limited to the instantaneous value of the agent's finite-range sensors. We endow the agent with an analog memory in the form of a linear time-invariant filter to provide the machine with a view of past sensory stimuli:

$$C_{1a} : \begin{cases} \dot{\boldsymbol{\xi}}_1 = \boldsymbol{A}\boldsymbol{\xi}_1 + \boldsymbol{B}(\boldsymbol{\phi}_1 - \boldsymbol{\delta}_1) \\ \boldsymbol{y}_{1a} = \boldsymbol{\xi}_1 \end{cases} \quad (13)$$

where $\boldsymbol{\xi}_1 \in \mathcal{R}^n$ and $n \in \mathcal{Z}$, $1 \le n < \infty$. The filter provides a crude model of the past, and its state, $\boldsymbol{\xi}_1$, provides information that the agent can use in controlling its circumstances. We place the memory block, $C_{1a}$, in cascade with block $C_{1b}$ (which realizes the control law), as shown in Figure 5.

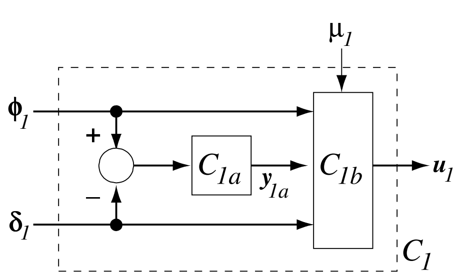


Fig. 5. Cascade of $C_{1a}$ and $C_{1b}$ to form $C_1$. The filter $C_{1a}$ serves as an analog memory.

We define the functions $\text{pul}_{l_1} : \mathcal{R} \mapsto \mathcal{R}$ and $\text{sat}_{l_1} : \mathcal{R} \mapsto \mathcal{R}$, as shown in Figure 6. The notation $\text{diag}(\boldsymbol{x})$ denotes a diagonal matrix in which entry $(i,i)$ is given by the $i$-th element of $\boldsymbol{x}$. The set $\Lambda_r$ is a half-closed interval of the form $[\varepsilon_r, \infty)$, where $\varepsilon_r > 0$, $\Lambda_{\theta,2}$ is any closed interval of the form $[\varepsilon_{\theta,2},\ \pi - \varepsilon_{\theta,2}]$ where $0 < \varepsilon_{\theta,2} < \pi$, and $\Lambda_{\theta,3}$ is any closed interval of the form $[-\varepsilon_{\theta,3},\ \varepsilon_{\theta,3}]$ where $0 < \varepsilon_{\theta,3} < \frac{\pi}{2}$.

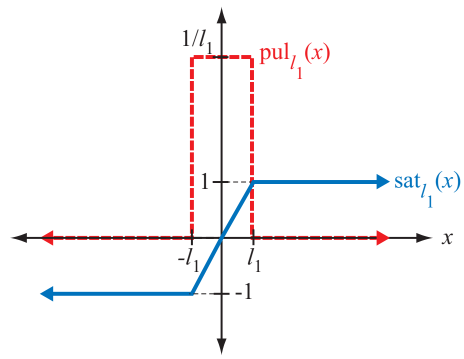


Fig. 6. Definitions of the $\text{pul}_{l_1} : \mathcal{R} \mapsto \mathcal{R}$ and $\text{sat}_{l_1} : \mathcal{R} \mapsto \mathcal{R}$ functions, parameterized by the constant $l_1 > 0$.

Due to the need for brevity, we present an overview of the results; detailed derivations can be found in [28].

Suppose that $\mu_1 = 0$, that is, the system is in target-seeking mode. The first result provides a control law that is valid in a domain $\mathcal{D}_0 := [\mathcal{R}^n \times (\Lambda_r \times \Lambda_{\theta,2} \times \Lambda_{\theta,3})] \subset \mathcal{R}^{n+3}$ and provides a simple basis for obstacle avoidance and taxic behavior.

*Theorem 1 (Stability on $\mathcal{D}_0$):* Given $\boldsymbol{\phi}_1 := \begin{bmatrix} \phi_r \\ \phi_{\theta,2} \\ \phi_{\theta,3} \end{bmatrix}$ where $\phi_r, \phi_{\theta,2}, \phi_{\theta,3}$ are real constants such that $\phi_r \in \Lambda_r$, $\phi_{\theta,2} \in \Lambda_{\theta,2}$ and $\phi_{\theta,3} \in \Lambda_{\theta,3}$, if $C_{1b}$ is given as:

$$C_{1b} : \boldsymbol{u}_1 = \boldsymbol{u}_{1,1} := \boldsymbol{B}_1^{-1}(\boldsymbol{\eta}_1) \left[ \boldsymbol{f}(\boldsymbol{\xi}_1, \boldsymbol{\phi}_1, \boldsymbol{\delta}_1) + \boldsymbol{B}^T \boldsymbol{P} \boldsymbol{\xi}_1 - \frac{\kappa_{1,2}}{2} \boldsymbol{\varepsilon}_1 \right] \quad (14)$$

where $\boldsymbol{P} \in \mathcal{R}^{n \times n}$ is a symmetric positive definite matrix,

$$\boldsymbol{f}(\boldsymbol{\xi}, \boldsymbol{\phi}, \boldsymbol{\delta}) := \left[ \boldsymbol{K}_{1,0} + \tfrac{\kappa_{1,1}}{2} \boldsymbol{\Gamma}(\boldsymbol{B}^T \boldsymbol{P} \boldsymbol{\xi}) \boldsymbol{B}^T \boldsymbol{P} \right] [\boldsymbol{A}\boldsymbol{\xi} + \boldsymbol{B}(\boldsymbol{\phi} - \boldsymbol{\delta})], \quad (15)$$

$\kappa_{1,1}, \kappa_{1,2} > 0$, $\boldsymbol{K}_{1,0}$ is chosen such that $(\boldsymbol{A} - \boldsymbol{B}\boldsymbol{K}_{1,0})$ is Hurwitz, $\boldsymbol{\Gamma} : \mathcal{R}^m \to \mathcal{R}^{3 \times 3}$ is given by

$$\boldsymbol{\Gamma}(\boldsymbol{x}) := \text{diag}\left( \text{pul}_{l_1}(x_1), \text{pul}_{l_1}(x_2), \text{pul}_{l_1}(x_3) \right), \ \boldsymbol{x} \in \mathcal{R}^3, \quad (16)$$

$\boldsymbol{\varepsilon}_1 := \boldsymbol{\delta}_1 - \boldsymbol{\delta}_1'$, where

$$\boldsymbol{\delta}_1' := \boldsymbol{K}_{1,0} \boldsymbol{\xi}_1 + \boldsymbol{\phi}_1 + \frac{\kappa_{1,1}}{2} \boldsymbol{\mathcal{S}}(\boldsymbol{B}^T \boldsymbol{P} \boldsymbol{\xi}_1), \quad (17)$$

and $\boldsymbol{\mathcal{S}} : \mathcal{R}^3 \to \mathcal{R}^3$ is given by

$$\boldsymbol{\mathcal{S}}(\boldsymbol{x}) := [\text{sat}_{l_1}(x_1) \ \ \text{sat}_{l_1}(x_2) \ \ \text{sat}_{l_1}(x_3)]^T, \ \boldsymbol{x} \in \mathcal{R}^3, \quad (18)$$

then, assuming the state vector $\begin{bmatrix} \boldsymbol{\xi}_1(t) \\ \boldsymbol{\eta}_1(t) \end{bmatrix} \in \mathcal{D}_0$ for all $t$, we have 1) $\left\| \begin{bmatrix} \boldsymbol{\xi}_1(t) \\ \boldsymbol{\eta}_1(t) \end{bmatrix} \right\|$ is bounded, and 2)

$\lim_{t\to\infty} \|\boldsymbol{\eta}_1(t) - \boldsymbol{\phi}_1(t)\| < \varepsilon_{l_1,\kappa_2}$, where $\varepsilon_{l_1,\kappa_2} > 0$ satisfies $\lim_{\substack{l_1 \to 0,\\ \kappa_{1,2} \to \infty}} \varepsilon_{l_1,\kappa_2} = 0$.

An overview of the proof is given in [28], but we note that the law is obtained by first invoking the Lyapunov function:

$$V_1 := \boldsymbol{\xi}_1^T \boldsymbol{P} \boldsymbol{\xi}_1 + \boldsymbol{\varepsilon}_1^T \boldsymbol{\varepsilon}_1 \tag{19}$$

and then applying Lyapunov-based synthesis, "high-gain" and backstepping techniques to derive a control law, $\boldsymbol{u}_1$, that ensures $\dot{V}_1 < 0$ for all $\boldsymbol{\xi}_1 \in \mathcal{R}^n$ and for all $\boldsymbol{\varepsilon}_1 \in \mathcal{R}^3$.

A unique aspect of our control strategy is that we intentionally relax the pursuit of asymptotic stability to ultimate boundedness [17] in order to improve obstacle avoidance behavior. Qualitatively, this provides some latitude enabling "counter-intuitive" decision-making to allow the robot to better cope with obstacles (e.g., backing up even though it means temporarily moving away from the target). We have found that as we approach asymptotic stability (through variation of $\kappa_{1,1}$), the robot becomes prone to being trapped, since it is less likely to deviate from an aggressive path towards its objective. The following result quantifies this feature.

*Corollary 1 (Intentional Relaxing of Control):* Let $\boldsymbol{R}$ be a symmetric positive definite matrix. If $l_1 \leq \frac{\rho_0 \|\boldsymbol{B}^T \boldsymbol{P}\|_1}{2}$ (where $\rho_0 := \frac{\sqrt{n}|\kappa_{1,1}| \|\boldsymbol{B}^T \boldsymbol{P}\|_1}{\lambda_{\min}(\boldsymbol{R})}$ and $\lambda_{\min}(\boldsymbol{R})$ is the smallest eigenvalue of $\boldsymbol{R}$), then setting $\kappa_{1,1} < 0$ induces a neighborhood of radius $\rho_0 > 0$ outside of which $\dot{V}_1 < 0$.

Now consider the case where $\mu_1 = 1$, and the system is in search mode. In this case, we generate the reference $\boldsymbol{\phi}_1$ from an internal pattern generator (a set of oscillators) and regulate the inertial guidance sensor, $\boldsymbol{\sigma}_2$. This will cause the agent to execute an approximation of a space-filling curve in the environment, that is, it will search the environment. From (12), we see that the velocity sensor output is simply $\boldsymbol{\sigma}_2 = \begin{bmatrix} v \\ \omega_2 \\ \omega_3 \end{bmatrix}$, in other words, no backstepping is needed, and the control $\boldsymbol{u}_1 = \boldsymbol{\delta'}_1$ is adequate (for slowly varying $\boldsymbol{\phi}_1$) to provide obstacle avoidance and velocity tracking.

*1) Comment:* An interesting side-effect of our synthesis is the emergence of normalized radial basis functions in the $\boldsymbol{\mathcal{S}}(\cdot)$ operator, and unnormalized radial basis functions in the $\boldsymbol{\Gamma}(\cdot)$ operator. Hence, $\boldsymbol{u}_1$ is produced, partially, by a radial basis function network.

### C. Level Two

The level two controller addresses the case when the robot is either:

1) in a region of space where the target is out of sensor range (i.e., $\sigma_3 = 1$)
2) trapped by an obstacle formation that the level one controller can not suppress (e.g., a concave formation)

We first construct a sense, $\delta_2$, that detects these problematic cases. For case one, we know $\sigma_3 = 1$. For case two, we note that if the level one controller is unable to suppress the obstacle, the integral of the obstacle sensor's radial component ($\sigma_{0,1}$) with respect to time will tend to increase (due to persistent stimulation of the obstacle sensor). Hence, we can filter the magnitude of $\sigma_{0,1}$ by a "leaky" integrator:

$$\dot{\zeta}_2 = -\kappa_2 \zeta_2 + |\sigma_{0,1}| \tag{20}$$

and apply a hysteresis function, $h_2 : \mathcal{R} \mapsto \mathcal{R}$, that triggers "high" (goes from 0 to 1) if $\zeta_2$ is increasing and crosses a threshold $T_{2,H}$, and resets (goes from 1 to 0) only if $\zeta_2$ is decreasing and crosses a threshold, $T_{2,L} < T_{2,H}$. We can now define:

$$\delta_2 := \sigma_3 + h_2(\zeta_2) \in [0, 2] \tag{21}$$

with $\delta_2 = 0$ being the desired state of affairs (i.e., the absence of the two problem conditions).

Define the actuation signal, $u_2 \in \{0, 1\}$, such that it controls the $\mu_1$ input of the level one controller, recalling that $\mu_1 = 0$ enables taxis, while $\mu_1 = 1$ enables searching. Now, we construct a plant that models the dynamics of $\delta_2$ in terms of $u_2$. Although we lack the information to construct a precise model (the resource-bounded agent lacks a global map and is constrained by sensors that can only produce local information), we can construct one that reflects the *trend* of how we *expect* $\delta_2$ to evolve with respect to $u_2$. Qualitatively, when $\delta_2 > 0$, we expect that engaging in searching (i.e., $u_2 = 1$) should tend to decrease $\delta_2$, whereas when $\delta_2 = 0$, $u_2$ should cause no change in $\delta_2$. We can describe this dynamically by:

$$P_2 : \begin{cases} \dot{\eta}_2 = -u_2 \\ \delta_2 = \eta_2 \end{cases} \tag{22}$$

(we note that by (21) $\delta_2 \in [0, 2]$, hence when $\delta_2 = 0$ and $u_2 \neq 0$ the model tends to bring $\delta_2$ out of its domain of definition; we thus need to verify that the final control law is consistent with the domain of definition of $\delta_2$).

To synthesize the controller that will cause $\delta_2$ to track $\phi_2 = 0$, we define the tracking error:

$$e_2 := \phi_2 - \delta_2 = -\delta_2 \in [-2, 0] \tag{23}$$

and the scaler-valued function:

$$V_2 := \frac{1}{2} e_2^2 \tag{24}$$

which is positive-definite with respect to $e_2$. Differentiating (24) with respect to time, we obtain:

$$\dot{V}_2 = e_2 u_2 \tag{25}$$

Setting $u_2 = -\text{sgn}(e_2) = \text{sgn}(\delta_2)$ makes (25) negative-definite with respect to $e_2$, and hence $\lim_{t\to\infty} \delta_2 \to \boldsymbol{\phi}_2 = 0$. We note that $u_2 \in \{0, 1\}$, since by (23) $e_2 \in [-2, 0]$—which is consistent with our earlier definition of the range of $u_2$.

*1) Comment:* The design of the level two regulator suggests how the methodology can be expanded to higher-order cognition. We note that here our plant model is derived from the qualitative trends of how we expect the world to react rather than a physically- and/or geometrically-precise derivation of how the world actually reacts. We feel that this way of divorcing ourselves from lower-level considerations is a step in the direction of synthesizing more abstract cognitive skills.

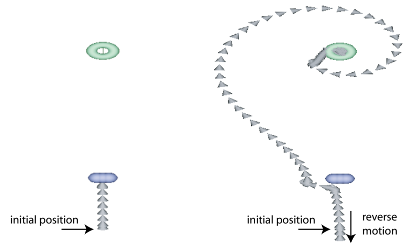


(a) $\kappa_{1,1} \geq 0$ (b) $\kappa_{1,1} < 0$

Fig. 7. Simulation #1: traces of robot trajectories for two cases illustrating the emergence of satisficing behavior as $\kappa_{1,1}$ goes negative.

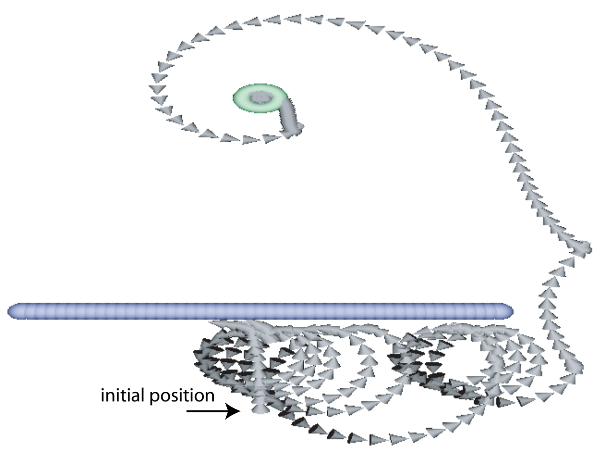


Fig. 8. Simulation/Animation #2: trace of robot trajectory; overstimulation by a large obstacle.

## V. Simulation

For clarity of presentation, we show an instantiation of the controller for the planar case[3]. The robot was placed at random initial conditions and simulated in various obstacle courses. Figure 7 illustrates the inability of a system with $\kappa_{1,1}$ to get around a small wall; when stability is relaxed by setting $\kappa_{1,1} < 0$ the situation is improved. Figure 8 illustrates the "searching" behavior elicited by the level two controller. Figures 9-10 illustrate the agent's behavior in more complex environments.

Animation videos for additional simulations are available from the following URL: http://www.ece.tamu.edu/~takis/robotics_main_page.html.

In the formulation of Section II the agent is dimensionless, whereas in our plots and animations a cone-shaped cursor was used to track the agent's location. This has resulted in the occasional appearance that the agent is simply bouncing off of obstacles. We remark that our simulation environment does not model any physics apart from the response due to acceleration commands from the agent—it does not model the physics of object-obstacle interactions. That is, in the absence of sensory feedback from $\boldsymbol{\sigma}_0$ regarding the presence of an obstacle, the agent in the simulation environment will move through obstacles as if they were not present. Hence, the observed deviation of the agent from obstacles is purely due to the agent's regulation of its sensory perception of the world.

[3]For the simulations in this paper, $C_{1a}$ was composed of two third-order elliptic low-pass 0.1 Hz filters.

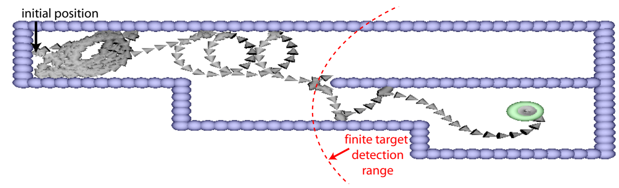


Fig. 9. Simulation/Animation #3: trace of robot trajectory; maze example with target initially out of sensor range.

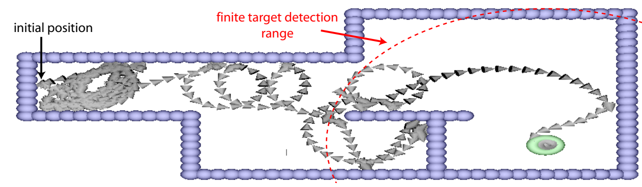


Fig. 10. Simulation/Animation #4: trace of robot trajectory; maze example with target initially out of sensor range.

## VI. Weak Emergence of Satisficing Intelligence

Simon [18], [19] suggests that cognitive systems are systems that *satisfice*, that is, systems that find "tolerable" rather than optimal solutions. Pollack [20] in her AI formulation of an intelligent agent introduces mechanisms that enable satisficing intelligence. With respect to our system, the behavior of Figure 7 is an example of satisficing intelligence—the agent takes locally non-optimal actions that allow it to get around the wall.

In [21], Bedau introduces the concept of *weak emergence*, reporting that weak emergence is manifest in all complex systems with [29] placing it as requisite property of complex adaptive systems. Bedau's definition states that a phenomenon, $P$, of a dynamical system, $S$, with dynamics specified by $D$, is *weakly emergent* if and only if $P$ can be derived from $D$ and the external conditions of $S$ but only by simulation.

We appeal to this definition to show how satisficing intelligence ($P$) is a weakly emergent property of the hierarchical dynamic system ($\{S, D\}$) specified in Sections III-IV.

*a) $\Rightarrow$:* We first show that the manifestation of $P$ is rooted in $\{S, D\}$. This is straightforward as through experimentation we observe that $P$ arises when $\kappa_{1,1} < 0$ (e.g., in the simulations of Figure 7). With $\kappa_{1,1} < 0$, the agent exhibits more varied behavior (including taking locally non-optimal actions) when it meets an obstacle and reverses and/or turns to circumnavigate the obstacle. When $\kappa_{1,1} \geq 0$, the agent is aggressive as it tracks the target—however, this *fanaticism* (to use the AI-inspired terminology of [30], [31]) prevents it from taking non-optimal deviations from its optimizing path towards the target.

Hence, the cause of $P$ can be traced to $\{S, D\}$ via the parameter $\kappa_{1,1}$—the degree to which the stability of $C_1$ is relaxed.

*b)* $\Leftarrow$*:* Now we show that $P$ can *only* be seen to emerge through simulation, that is, we can not derive its manifestation purely by analyzing $\{S, D\}$. We note that in the synthesis of Section IV, we did not explicitly design behavior for turning around obstacles. This was because the agent we designed only had access to *local non-directional* obstacle sensing. Hence, we simply did not have access to the information required to design a regulator that could directly maneuver around an obstacle. Thus, by solely considering $\{S, D\}$ it is not possible to deduce the emergence of $P$ because the regulators in $\{S, D\}$ do not receive sufficient information to, by design, engage in $P$-like behavior.

What we did design into the system through Corollary 1 was a relaxed requirement for stability. That is, we lessened the *constraint* on the level one controller providing it with the *freedom* to take more varied actions—but we can not say what it will exactly do. Interaction with the environment—through simulation—is necessary to observe the manifestation of $P$.

### A. Comments

We have argued that the emergence of satisfying intelligence satisfies Bedau's definition of weak emergence. Moreover, we have shown how this is rooted in a unique feature of our architecture: the relaxation of stability of Corollary 1.

Traditional control-theoretic design methodology pursues aggressive optimization, i.e., asymptotic stability, with ultimate boundedness being the "next best thing" to be sought only when asymptotic stability is not possible. As we have seen, however, while asymptotic stability may lead to optimality, it does not necessarily lead to intelligence. We feel this insight is important to "tame" the application of control-theoretic tools so that cognition for "life-like" artificial agents can be synthesized.

## VII. Conclusion

Recognizing:

- the vital role of cognitive processes in life
- the prevalence of hierarchical dynamic systems in living processes and cognition
- the importance of embodiment in cognition

we sought to develop a cognitive architecture for an embodied robotic agent. Specifically, we composed an embodied hierarchical system by replicating regulator motifs, utilizing a methodology of plant-controller co-design to approach this synthesis rigorously. We note that despite the mathematical sophistication of the toolsets used, the resulting control system is amenable to economical analog implementation; data converters—expensive requirements of software-based control—are not needed.

Beyond the simulation results and animations which illustrate the functioning system, we point to the weak emergence of satisficing intelligence as a demonstration of the utility of our methodology in synthesizing the cognitive faculties for artificial life.

## References


[1] G. Sommerhoff, *Analytical Biology*. Oxford University Press, 1950.

[2] W. R. Ashby, *An Introduction to Cybernetics*. Chapman and Hall, 1968.

[3] H. R. Maturana and F. J. Varela, *Autopoiesis and Cognition*, ser. Boston Series in the Philosophy of Science. D. Reidel Publishing Company, 1980, vol. 42.

[4] R. A. Brooks and L. Steels, Eds., *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*. Lawrence Erlbaum Associates, 1995.

[5] W. R. Ashby, *Design for a Brain*, 2nd ed. Chapman and Hall, 1960.

[6] V. Braitenberg, "Taxis, Kinesis and Decussation," *Progress in Brain Research*, pp. 210–222, 1965.

[7] ——, *Vehicles: Experiments in Synthetic Psychology*. The MIT Press, 1984.

[8] M. A. Arbib, *The Metaphorical Brain*. Wiley-Interscience, 1972.

[9] J. S. Albus, "Mechanisms of Planning and Problem Solving in the Brain," *Mathematical Biosciences*, vol. 45, pp. 247–293, 1979.

[10] R. A. Brooks, *Cambrian Intelligence: The Early History of the New AI*. The MIT Press, 1999.

[11] M. J. Matarić, "Studying the Role of Embodiment in Cognition," *Cybernetics and Systems*, vol. 28, pp. 457–470, 1997.

[12] R. D. Beer, *Intelligence as Adaptive Behavior*. Academic Press, Inc., 1990.

[13] T. Ziemke, "What's that Thing Called Embodiment?" in *Proceedings of the 25th Annual Meeting of the Cognitive Science Society*, 2003.

[14] R. D. Beer, "A dynamical systems perspective on agent-environment interaction," *Artificial Intelligence*, vol. 72, pp. 173–215, 1995.

[15] V. J. Lumelsky, "Algorithmic issues of sensor-based robot motion planning," in *Proceedings of the 26th IEEE Conference on Decision and Control*, vol. 3, 1987, pp. 1796–1801.

[16] C. Mead, *Analog VLSI and Neural Systems*. Addison-Wesley, 1989.

[17] M. J. Corless and G. Leitmann, "Continuous state feedback guaranteeing uniform ultimate boundedness for uncertain dynamic systems," *IEEE Transactions on Automatic Control*, vol. 26, no. 5, pp. 1139–1144, October 1981.

[18] H. A. Simon, "A behavioral model of rational choice," *The Quarterly Journal of Economics*, vol. 69, no. 1, pp. 99–118, 1955.

[19] ——, "Cognitive science: The newest science of the artificial," *Cognitive Science*, vol. 4, pp. 33–46, 1980.

[20] M. E. Pollack, "The uses of plans," *Artificial Intelligence*, vol. 57, pp. 43–68, 1992.

[21] M. A. Bedau, "Weak Emergence," in *Philosophical Perspectives: Mind, Causation, and World*, J. Tomberlin, Ed., 1997, vol. 11, pp. 375–399.

[22] M. A. Bedau, J. S. McCaskill, N. H. Packard, S. Rasmussen, C. Adami, D. G. Green, T. Ikegami, K. Kaneko, and T. S. Ray, "Open Problems in Artificial Life," *Artificial Life*, vol. 6, pp. 363–376, 2000.

[23] S. Rasmussen, N. A. Baas, B. Mayer, M. Nilsson, and M. W. Olesen, "Ansatz for Dynamical Hierarchies," *Artificial Life*, no. 7, pp. 329–353, 2001.

[24] M. N. Jacobi, "Hierarchical Organization in Smooth Dynamical Systems," *Artificial Life*, vol. 11, pp. 493–512, 2005.

[25] S. McGregor and C. Fernando, "Levels of Description: A Novel Approach to Dynamical Hierarchies," *Artificial Life*, vol. 11, pp. 459–472, 2005.

[26] M. Krstić, I. Kanellakopoulos, and P. V. Kokotović, *Nonlinear and Adaptive Control Design*. Wiley, 1995.

[27] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Prentice Hall, 2002.

[28] T. Zourntos and N. J. Mathai, "A BEAM-Inspired Lyapunov-Based Strategy for Obstacle Avoidance and Target-Seeking," in *American Control Conference*, 2007 (to appear).

[29] J. H. Holland, *Adaptation in Natural and Artificial Systems*, 2nd ed. MIT Press, 1992.

[30] M. Bratman, "Two faces of intention," *The Philosophical Review*, vol. 93, no. 3, pp. 375–405, July 1984.

[31] P. R. Cohen and H. J. Levesque, "Intention is choice with commitment," *Artificial Intelligence*, vol. 42, pp. 213–261, 1990.