

LIVING MACHINES



Figure A: “Bug of the badlands”

Brosi Hasslacher
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
<bhass@lanl.gov>

Mark W. Tilden
Physics Division
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
<mwtilde@lanl.gov>

Living Machines

Brosi Hasslacher
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
<bhass@lanl.gov>

Mark W. Tilden
Physics Division
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
<mwtilde@lanl.gov>

Abstract

Our aim is to sketch the boundaries of a parallel track in the evolution of robotic forms that is radically different from any previously attempted. To do this we will first describe the motivation for doing so and then the strategy for achieving it. Along the way, it will become clear that the machines we design and build are not robots in any traditional sense. They are not machines designed to perform a set of goal oriented tasks, or work, but rather to express modes of survivalist behavior: the survival of a mobile autonomous machine in an a priori unknown and possibly hostile environment. We use no notion of conventional "intelligence" in our designs, although we suspect some strange form of that may come later. Our topic is survival oriented machines, and it turns out that intelligence in any sophisticated form is unnecessary for this concept. For such machines, if life is provisionally defined as that which moves for its own purposes, then we are dealing with living machines and how to evolve them. We call these machines biomorphs (**BIO**logical **MORPH**ology), a form of parallel life.

Introduction to Biomorphic Machines

One difference between biological carbon based life forms and the mobile survival machines we will discuss are materials platforms, which can be metals, plastics, silicon -- a large variety of materials, but as we take these principles and descend in scale, it becomes clear that we could use many of the ready made protein structures provided by molecular cell biology for other purposes. We could in principle self-assemble this new machine life out of carbon chemistry at the scale of the cell, using the cells' ingredients, power sources and ATP engines. Doing that would not recover cell-centered biology, but manufacture novel, potentially useful life forms that have apparently escaped the normal path of evolution.

What is different about biomorphic machines from typical mobile platform designs is not their materials base but how they are organized. They use a dynamical, non-symbolic internal world representation and compliant, bi-directional, interactive response where the external world assumes a crucial role. In this they have much in common with biological forms which is not accidental; these machines are designed along biological paradigms rather than on first principle notions of how such machines should be organized. We take the viewpoint that such principles will have to be discovered by experiment rather than postulated by pure reason. In a sense these machines are

evolved by physicists and engineers looking through biological eyes. Ultimately however, biomorphs are self designed by the machines' own emergent survival capabilities.

This is a study in experimental machine morphology and psychology. Over 70 working machines have been built and principles extracted to design more efficient but less complex machines with better cost-verses-survival functionality. We consider this field an experimental science in which we both learn from the machines and are for the moment their evolutionary agents.

The simplest way to describe a biomorphic device is to say that the whole machine acts as an analog computer, designed along biological paradigms, to move in, interact with, and survive in an unknown but fractal external world. There is no notion of programming, but rather adaptive, parallel reconfiguring of signals in neuron circuits, typically in ring topologies. These structures compute, but not in any digital sense. This leads to the idea of a biomorphic architecture.

Biomorphic Architectures and Global Machines

Biomorphic architecture is autonomous machine architecture modeled on compliant skeletal wandering mechanisms found in biology. All effective computation is done in analog and from the periphery inward. It is modular and tiered. No digital computation is ever done within the motion platform, although digital pulse trains are used for motor drive and control. The essence is a core of electronic neurons that is bidirectionally connected to standard sensors and "smart" mechanical appendages that locally do much of the immediate computation necessary for their function.

From a systems viewpoint the entire mechanism is a single analog computer with a local modular architecture, where analog computation occurs in two realizations: mechanical at the machine periphery and electronic at the core. The simplest non-trivial morphic architecture is a mobile survival platform using only mechanics and an electronic core. This allows for sufficient, if not efficient, negotiation of undefined, complex terrains.

Operationally, a biomorphic machine is a global rather than local object despite its modular construction. The entire device acts as a unit with mesoscale properties that could not be inferred from a description of its components. In this way it resembles cellular automata which are collections of simple finite state machines whose time evolution proceeds by simple rules. Nevertheless the range of behaviors possible from such a setup extends from trivial fixed point behavior to Turing machines. These behaviors are collective or emergent and arise at the mesoscale of the system. An example is the cellular automata model for the Navier-Stokes equation for fluids; the collective behavior at the mesoscale of a very simple set of cellular automata rules. In biomorphic architectures we do not have finite state machines in the normal sense, since there is no concept of digital or symbol, nor preset update rules, only dynamic interactions among the parts of the machine. However the analogy is familiar territory and a good one to keep in mind.

Although trivial pieces of biomorphic machines, such as an entire leg can be removed, altered or damaged without altering the machines' behavior, modifications to the internal neural architecture will alter global response drastically. Most biomorph machine "nets" employed so far are in loop and/or link ring structures with a single minimal control core that once set cannot be topologically altered without creating a completely

new class of behaviors. There is no concept of smooth deformation away from the innermost minimal core. The way to increase the functionality of the machine is to add additional functional ring structures that can be smoothly deformed. The upper bound for such designs is obviously infinite, but there is an efficient lower bound architecture that implies the idea of a minimal survival neuron network that we will describe, following some background theory on general living machine architectures.

Soft Machines

First we describe a setup designed to produce sharp mental images with a minimum of formalism and introduce a lexicon that we have found convenient to describe these machines.

A soft machine is a biologically based concept in which the machine forms its behavior through interactions with a complex and a priori unknown environment. This changes with time and interaction; it is dynamic, and has few, if any state consistencies between modes. It is put in opposition to a hard machine whose behavior is sharply defined from the start by look-up tables, branching logic or other conventional programming schemes.

Walkers

One imagines the following picture: there exists some roughly fractal world that we wish a machine to negotiate and survive. The machine should extract power from its environment -- power can be thought of as a form of food and it should always look for better and more reliable sources of power, but this ideal case, however desirable and practical, is not necessary to the idea. In practice the fractal world is some mechanical terrain and the machine must move and react to it. The terrain is always assumed dynamic and even hostile -- populated with other life forms moving for their own reasons and searching for food. We choose a walking mode rather than flying, swimming or wheels as a first step for several reasons. First a walker can efficiently negotiate severe terrains without constriction, a lesson learned from biology. Second, walkers provide very important visual clues to their builders as to whether they are operating properly and aids immeasurably in indicating solutions to the convergence of the "creatures" neural core. Because most higher biological life forms are walkers, people have evolved an acute sense of "body language" allowing us to immediately recognize whether a walker is operating properly and in what mode. With wheeled devices this is almost impossible. Again we are taking many cues and strategies from biological behavior.

Layered Autonomy

Next we use an important clue from biological architectures, the idea of layered autonomy, around which the entire architecture is constructed.

Mechanical Layer

Biomorphic walker legs must be able to solve the problems of balance dynamically as independent units, separate from knowledge of or help from the rest of the organism.

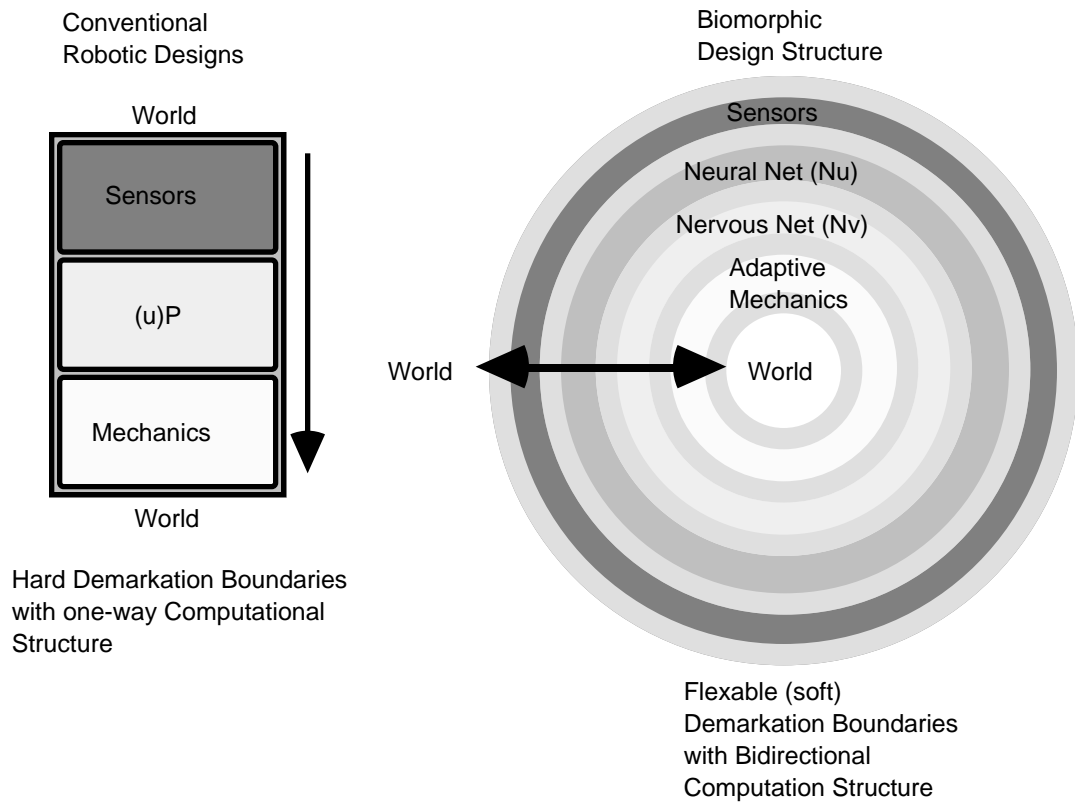
This can be done in several ways. We chose the simplest in conception but not in execution, namely "smart" mechanical legs with low degrees of freedom (DOF) but high degrees of structural compliance. They compute balance in mechanical analog, using biologically motivated mechanical structures with fuzzy logic joints and unavoidable stochastic slop. There is no digital processing whatsoever done on the machine. It is both unnecessary and, in implementation, very costly in terms of generalized computational resources. Each leg balances itself in parallel with the others. Compliant mechanical legs are also appropriate objects for interacting with a hard fractal world. They can be made of such size and materials to negotiate complex and rough terrain with minimal damage to either walker or environment. As well, they can double as manipulators so that the function/complexity ratio of the machine is optimized.

Neuron Core

In analogy to biological organisms, we use an artificial nervous system with adaptive control to produce appropriate adaptive walking gaits for these machines. We equip the legs with explicit and implicit local sensors that allow it to make a highly abstract image of its immediate environment. Distal sensors are irrelevant for sufficient minimalist machines; they do not need the cognitive skills to process such information. Explicit sensors can be of many kinds, for the purposes of a simple picture one can imagine them to be simple impact or proximity sensors. Implicit sensing in the biomorphic case is in the form of torque feedback from all leg motors, giving a direct and sufficient indication of terrain complexity. These signals are carried by the electronic analog of a nervous bundle to the central neuron core which shapes and delays a cycle of timing signals to the drive motors in a dynamic fashion. If the nervous core has been calibrated appropriately for the mechanical elements of its physical body, the emergent behaviors are recognizable walking gaits for the machine.

All of this processing is analog, constructing a primitive, dynamic, but sufficient internal neuron core representation of the machine's state in the external world. As it is the only representation biomorphs have of the external world, this sharply distinguishes it from strategies using highly detailed, preordained internal world representations and opposite strategies that use no such representation. Both these conventional strategies have proven, by experiment, to be extremely costly in computational resources. In contrast, we have capable walking machines that negotiate complex unknown terrain using a total of twelve transistors as the computational core.


Figure 1: Comparison of Conventional vs. Biomorphic Architectures



The output of the neuron core propagates down the neural bundle to drive motors and sensors, modifying their behavior and closing the loop. As the machine interacts with the external world, the internal representation of that world changes continuously. The neuron core acts as a variable-rate, short term memory whose independent components have no knowledge of belonging to a larger organism.

The overall design picture that emerges is a single global analog computer whose various pieces change realization. They smoothly morph through several stages: sensors that deliver modification information to a heuristic neural net "brain," which influences an independent neuron machine core that contains a highly abstract and condensed representation of the external world. Outwards the mechanical body, also computing continuously, can interact with the external world, but has no knowledge that it even belongs to an integrated creature: a sum of components, all with soft demarcation boundaries, blending to form survival adequacy. Figure 2 gives a typical design representation for a quadruped creature, which can be quickly extended to 6, 8, or $2n$ legged devices (odd number legged devices are possible but ineffective as the odd leg induces a drag on the structure that the control core can make little use of).

Figure 2: Biomorph Computational Structure



Internal Landscapes and Roaming Space

Biomorphic survival traits and behaviors can be emphasized, not set, by the careful matching of the variables among layered components. To represent this, we chose to equip biomorphic machines with an expandable internal landscape that self-assembles an internal abstract world. This internal representation resides primarily in the nervous-network core (Nv) and has severe constraints. Elementary walking gaits are functional blocks that become atomic structures. Through interaction with sensors and motor loads, they loosely couple the whole machine into a single system with a global world representation, giving the machine complex dynamical systems properties. One learns from dynamical systems theory never to couple the parts of a nonlinear machine too tightly if you want complex behavior. Tight coupling, as in programmed tasks, constrains the development of emergent nonlinear behavior to the point of extinction. These machines display emergent behavior; from a systems viewpoint the entire machine is a loosely coupled, parallel computer. This allows biomorph designs not just adaptation to complex terrain, but also resilience against element damage. Indeed, we have seen that such designs can withstand up to 80% damage of their systems yet still attempt to continue moving. At this point they are far from efficient, but such designs may be sufficient for many tasks. So long as a mechanism can remain moving, it is capable of continuing performance.

The neural core plays a special role in this picture, for it is here that the distillate of all computation done by the periphery of the machine is focused in an abstract and condensed

form. Here presentation is non-symbolic, consisting in the delay and shaping of a chase-series of timing pulses that drive the motors of the device. The world is a set of pulse shapes traversing a ring topology and constantly attempting phase lock synchronization among variable pulse trains. There is an internal /external balance set up that is flexible enough to give both reactive and emergent behaviors to the machine.

World Representation

There are two types of approaches to the world representation problem in current use. One uses no internal representation and the other uses a highly detailed representation, which though clever, immediately dooms the machine to paralysis or destruction in any environment too far from its fixed response list. A future aim for biomorphic designs is to see if a simple self-assembled and dynamic internal world representation is a viable alternative. The verdict on this is so far open, but biomorphic representations do appear capable of a sophisticated and strange type of emergent intelligence that we may not immediately recognize. By this we mean the machine develops the ability to find adaptive solutions to complex adaptive problems that are analog, purely parallel, and not symbol based. Intuitively, such designs would seem to be contradictory in their abilities. Experimental evidence shows, however, that these devices are not just passively convergent on survival solutions, but aggressively so.

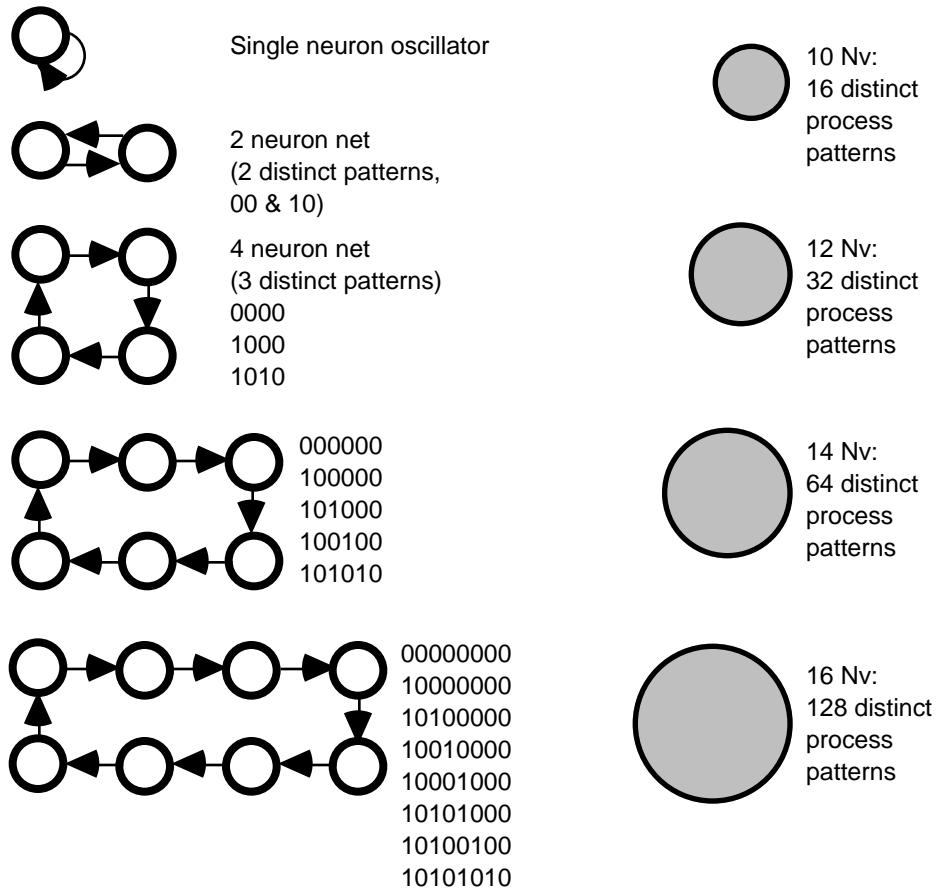
The ring structures of biomorph nervous net design hinges on the emergent computational properties of biological motor neurons as compared to the signal adaptation abilities of classic neural structures. Our artificial motor neurons work in topological chains, loops, and intersections and act as effective pre-processing elements between the motors and whatever controlling "head" may direct the mechanism through its environment. The design structure of these networks is not unlimited but is constrained by the dimensional limits of the machine's morphology, center of balance, power availability, and motor efficiency. The phase space is vast but if designers concentrate on minimalist arrangements, elegant, competent designs emerge.

Intelligence

A biomorphic architecture roams in the world and interacts with it, concurrently constructing a dynamic complex internal representation of the external world in its neuron landscape. We could enlarge their internal landscape in an initially unconstrained way, for example by providing neurons on an analog VLSI chip with free nodes that the machine could adaptively explore and connect. We conjecture that a form of sub-cognitive intelligence emerges by interacting and surviving in a hostile and unknown environment, provided the machine can encode its experience, not by symbols, but by altering dynamically the connection matrix of a parallel internal landscape. How large this landscape must be is still an experimental question, but early indications are that it can be surprisingly small.

A nervous net (Nv) acts as a medium that supports independent processes passed from neuron to neuron based upon independent neuron timing values. The range of this process space and their interaction dynamics are shown for some of the smaller, simple ring structures in figure 3:

Figure 3: Some Distinct Process Patterns of Artificial Motor Neuron (Nv)



Individual processes (1s) are independent of each other in time so long as they remain more than two neurons away from other processes. When processes get close, they mode lock themselves into a common, synchronized time base determined by the propagation time t of the trailing process at that neuron. The result is, using a 6 neuron core as an example, that in the two-process pattern 101000, both will cycle in synchronization forever, provided one neuron in the chain has a shorter t than all the others. When all neurons have approximately the same t , the two processes eventually mode lock into the 100100 pattern, where they loop in a fragile 180 degree mode lock. In a three-process loop, 101010, the entire process chain rotates at the speed defined by the shortest Nv node delay (called the "roller-coaster point"). What this means is that the more processes that are introduced into an unbalanced Nv loop, the faster it will travel. This implies that walking robots using this pattern will be able to increase not only the number of legs they can use for each gait, but also the speed at which those legs are used. Consequently, biomorphic robots have an inherent gait stabilization ability regardless of the size of their internal controller; if more legs are called into play, they "run" faster as a natural consequence of Nv process physics.

Complexity

One can give a rough argument that the high degree of complexity that results for such designs is possible and perhaps unavoidable. Suppose we had an analog VLSI chip populated with neuron nodes and arranged so that the machine had dynamic access to node connectivity. This means that at every time step the chip was so arranged that the machine could create and destroy links among various nodes in a connectivity matrix. This is different from conventional neural network configurations as all structure elements are time independent of each other.

This type of representation is independent of spatial embedding. It also means that a mechanism exists on the chip so that there are two types of distinct operations, creation operators that connect two nodes and annihilation operators that can cut or destroy the links between nodes.


The nodes themselves are not static. They contain a signal that varies in time, approximating conventional forms of cellular automaton rule updating. Similarly the creation and breaking of node links can be thought of as separate cellular automata updating according to a different set of rules.

The setup in skeletal form consists of two independent sets of cellular automata, the node set and the link-unlink set that update according to rules governed by environment-machine interactions. Now we loosely couple these two sets of automata according to a third set of coupling rules also set by the machine. Such a setup is known to produce behavior ranging from collapse of the grid to a fixed point, to exponential explosion of connectivity. It also has intermediate regimes that are capable of emulating the dynamics of strong hyperbolic dynamical chaotic systems. A fourth regime exists in which strong local fluctuations of link connectivity occur but the overall Hausdorff dimension of the system tends rapidly to stable values.

The result is a structure extremely rich in its ability to store and connect information.

Super Layers - Neural Nets

There are many traditional alternatives available that allow biomorphic machines to do useful work. If we take a biological paradigm and slave each nervous neuron from a conventional neural (cognitive) neuron, we now have a structure that is baseline adaptive, but with a heuristic, semi-cognitive shell. Sensor structures placed outside this structure are processed by the neural layer to give secondary stimuli to the nervous nets dynamic problem solving abilities. This was done in a biomorphic device called "Lobster".



VBUG 1.2 "LOBSTER"

Single battery, 0.9 Kg. Metal construction, exoskeletal framework.

6 tactile sensors; 2 antennae, 4 leg.

Control core:

12 transistor adaptive Nervous Net (Nv)

16 transistor heuristic Neural net (Nu)

40 transistor motor drive array.

Total: 68 transistors

Emergent behaviors:

-learns walking in 3 sec. from cold start.

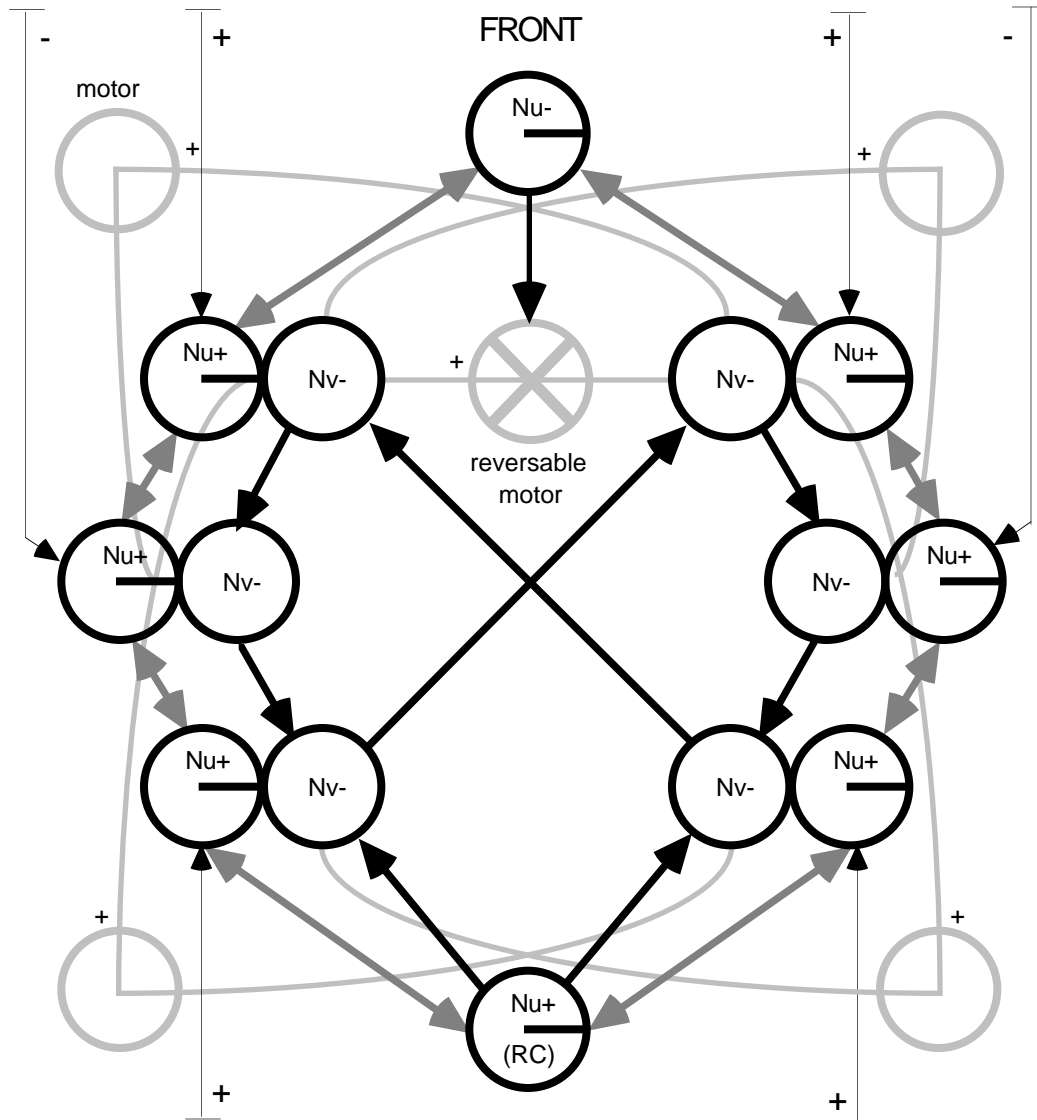
-obstacle avoidance, retreat, attack.

-Nu net accelerates Nv net learning ability by over twice, with 64 possible "moods".

-6 distinct walking gaits (stop, pace, trot, cantor, pronk, crabwalk)

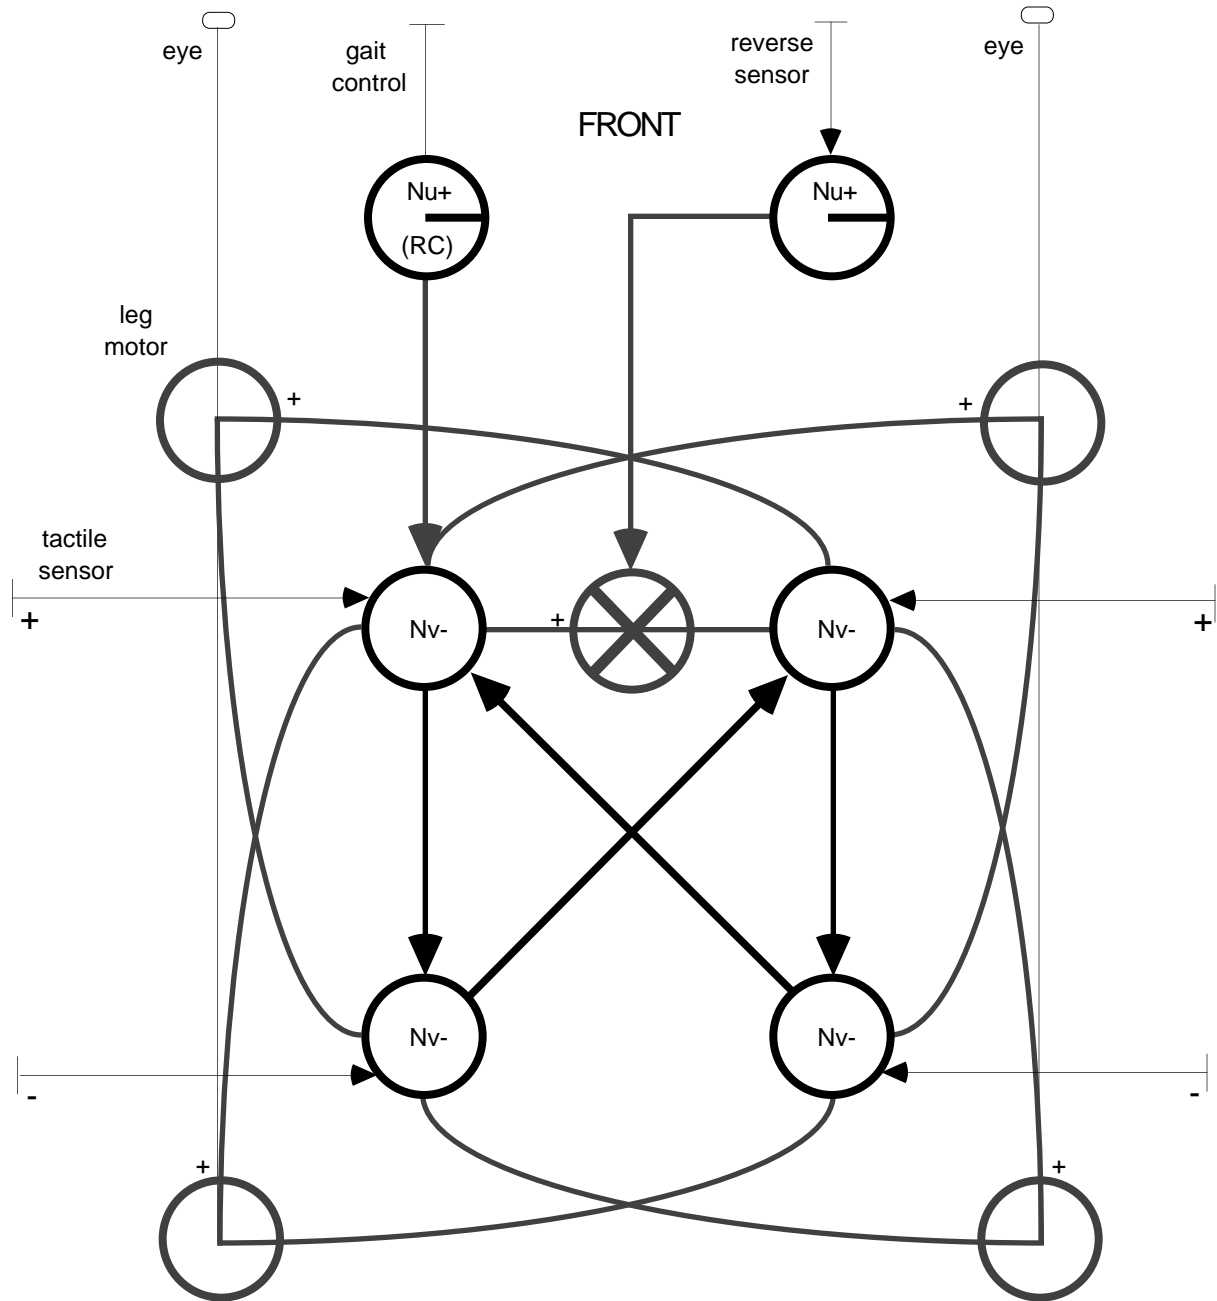
The complete nervous structure is shown in Figure 4. As seen, the device is almost a perfectly concentric biomorphic map. The neural layers mapped a total of 64 possible response influences (based upon external stimulus) onto the 6 distinct patterns the nervous ring was capable of sustaining. The device was quite capable in low stimulus environments but became "confused" when its environment became complex (i.e., dynamic, rather than stable external stimulus). In such cases it tended to increase gait cycles as an attempt to escape the stimulus, and if not possible, eventually locked itself in a catatonic condition where it remained until the environment settled. Baseline emergence revealed itself as the Nv cores ability to dynamically re-balance the machine in complex environments up to one-half the creature's height.

Figure 4: Vbug 1.2 "Lobster" Complete Neural Structure



Lobster taught us that combining neural (Nu) and nervous (Nv) neurons not only yielded much more elaborate behaviors than just neural structures alone, but that designs could be much denser if we assumed that the network electronics was well shielded against damage (damage that would otherwise force nets into lesser process pattern space). Since the Lobster neural structure slaved 2 neurons in 3 ways, and the remaining 4 only one way, it was possible to "row-reduce" the network further to a point where all neurons were driving at least two motor actions. Initially, this pattern seemed too interdependent to assure convergence, especially since reducing the nervous core to only 4 neurons would limit possible gaits to just 3 (stop, walk, dig) from the 6 currently possible. Surprisingly, rather than a low ability failure, "Walkman" as it is called, turned out to be the most capable of the 6 experimental walkers built to date. The complete neural structure for Walkman is detailed in Figure 5:

Figure 5: Vbug 1.5 "Walkman" Complete Neural "Microcore" Structure




This figure shows the smallest possible nervous network (defined as a "Microcore") for a capable quadruped with 1.25 DOF per leg. Unlike any larger ring designs, if sections of this net are damaged the results are fatal. The remaining network would not be rich enough to sustain a sufficient process spectrum that could generate a walking gait. The microcore is composed of three morphic elements: Sensors, Neural-net neurons (Nu) and the motor neurons (Nv). The Nu neurons first filter the activation signals in the Nv net to regulate gait processes, and secondly effect a change in the topological Nv structure to allow the robot a back-up ability. Without this regulation, the Nv 4-node core is sufficient to handle real-world processing, but not capable of regulating its

responses on power-on. "RC" refers to "Reticular Cortex" as an analogy to the biological mechanism found in living nervous systems that regulate excessive involuntary actions.

Because of the flexibility of the symbol lexicon used, the microcore diagram is not just a connection map but also an accurate position map of the robots' limbs and sensors. Looking at the robot from the top, sensors, motors and control core are all in their appropriate topological position. This is possible for such minimal designs but may fail when designs reach a greater complexity. Presently we find it very useful.

There were 3 major results from Walkman.



VBUG 1.5 "WALKMAN"

Single battery. 0.7Kg. metal/plastic construction. Unibody frame.

5 tactile, 2 visual sensors.

Control Core: 8 transistor Nv.

4 tran. Nu, 22 tran. motor.

Total: 32 transistors.

Behaviors:

- High speed walking convergence.
- powerful enviro. adaptive abilities
- strong, accurate phototaxis.
- 3 gaits; stop, walk, dig.
- backup/explore ability.

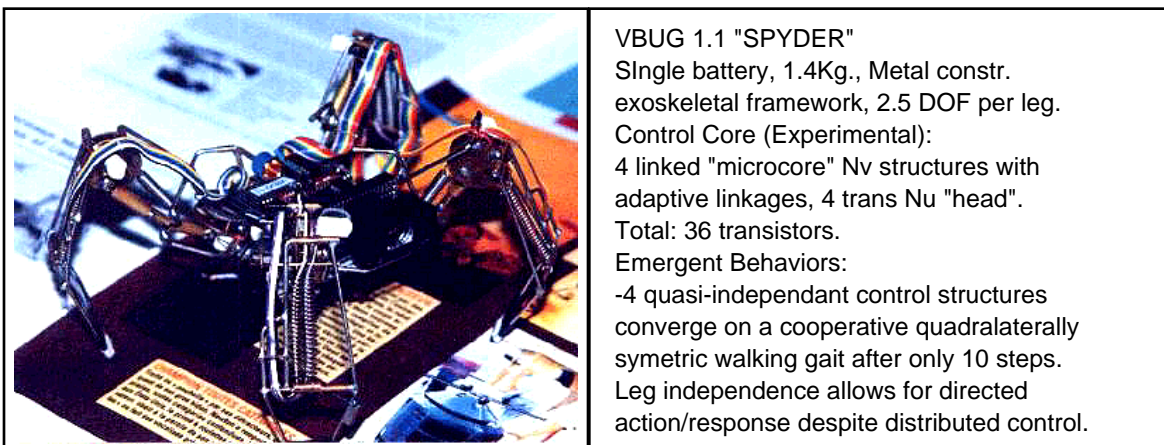
First, although the numbers of walking gaits were reduced, the network's ability to converge on a sufficient walking solution is very fast. The standard test is to twist all the legs on a biomorph machine 180 degrees out of walking phase and see how many steps the robot takes to achieve forward motion on a level surface. Biomorph legs, by convention, are completely unconstrained to take advantage of the largest possible effective area. Turtle, a basic, non-sensored Nv net on legs physically similar to Lobster (6 neuron core), sorts its legs out for forward locomotion in 14 steps. Lobster, because of its heuristic assist, manages in only 7 steps. Walkman takes only 1.5 steps and has been seen to frighten researchers and various domestic animals as a consequence.

Second, although there are only 4 time domain variables needed to converge a solution for this network, it took well over a day to find these values by experiment, as the degree of influence each neuron had on the others was enormously increased. Many sufficient solutions were found, but these often favored either drive or lift (the two essential walking operations) but not both. Considering the sum of biases on each neuron (Walkman was, like all other biomorph designs, far from physically symmetric) the final solution caused an exaggerated stepping gait that allows Walkman to climb obstacles twice its own height, and lower itself down from obstacles 4 times its height. Interestingly, this gait also made Walkman very difficult to high-center. By placing the device on a hockey-puck sized platform that suspended all feet from contact, it did manage to eventually gain a foothold and escape, by whipping its limbs around and using angular momentum to move its torso upon the platform. It was an interesting thing to

watch, and something a wheeled device can not manage, as anyone who's had their car trapped on top of a snow drift can attest.

Third, the terrain handling abilities of Walkman are unexpectedly vast. The device was equipped with pointed rubber feet to give it high traction on smooth surfaces and observed many times as it tried to figure its way across a cluttered, equipment filled desk. Even against the formidable task of having to crawl over stacked coat hangers, Walkman eventually found a solution. However, there was one drawback in the counterintuitive aspect of Walkmans' success; survival ability had gone up exponentially with a linear decrease in device complexity. This was annoying as it implied that a much more complex creature, the "Spyder" walker then under development, would not have anywhere near the survival metric of this far simpler creature.

The Microcore points to a whole new way of looking at creature designs. If we assume the microcore as a single, computing element, could clusters of these cores produce much larger internal computational spaces? Could overall survival behaviors emerge from a microcore cluster? As the smaller robots were of fixed designs, this was attempted on the much larger, mechanically complex Spyder, shown below.

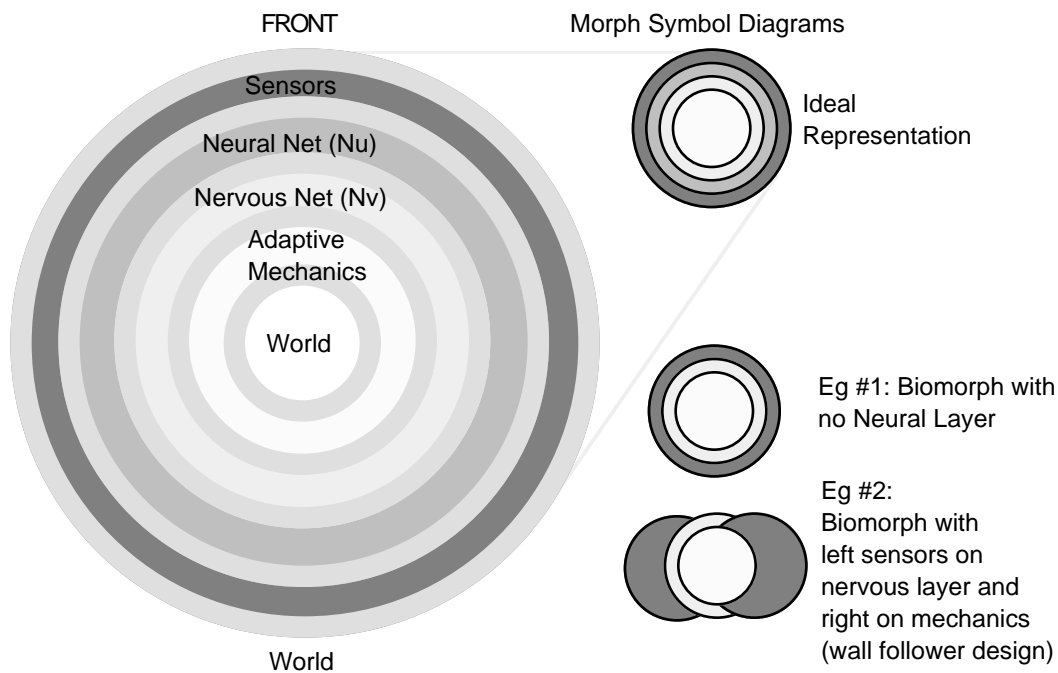


An original attempt at hand-converging a tiered Nv core for Spyder took over two weeks to arrive at a sufficient walking gait, because of the splayed mechanical dependence of 8 motors on a high-compliance, 4-bar linkage frame. Surprisingly, using the microcore-cluster controller, Spyder exhibited a sufficient walking gait on first power on, and an efficient gait only a few adjustments later. The implications are that microcore clusters can act as adept local processors that are capable at this higher design plateau. It also implies, as each cluster represented one leg in this design, that social biomorphism was a distinct possibility even at primitive levels of complexity. Micro-cluster Spyder could be theoretically carved up into permutations of quadrilateral slices and each permutation could still be expected to exhibit emergent, directed motion. Because Spyder was not originally designed for vivisection, this theory has yet to be proved in actuality, but a new generation of cooperative microcore slices is being designed.

The design space for these minimalist devices has, with this observation, expanded exponentially. We can imagine that groups of biomorphic robots could aggregate in social loops, chains, or even three-dimensional hives with a good chance of recursive symbiontism. That is, we can now build a "hive" group of minimalist microcore robots and assume this hive will definitely have a larger collective survival space than the individual spaces of its cellular parts. This will not be confirmed until such a hive is made operational and observed. This is a topic of ongoing research.

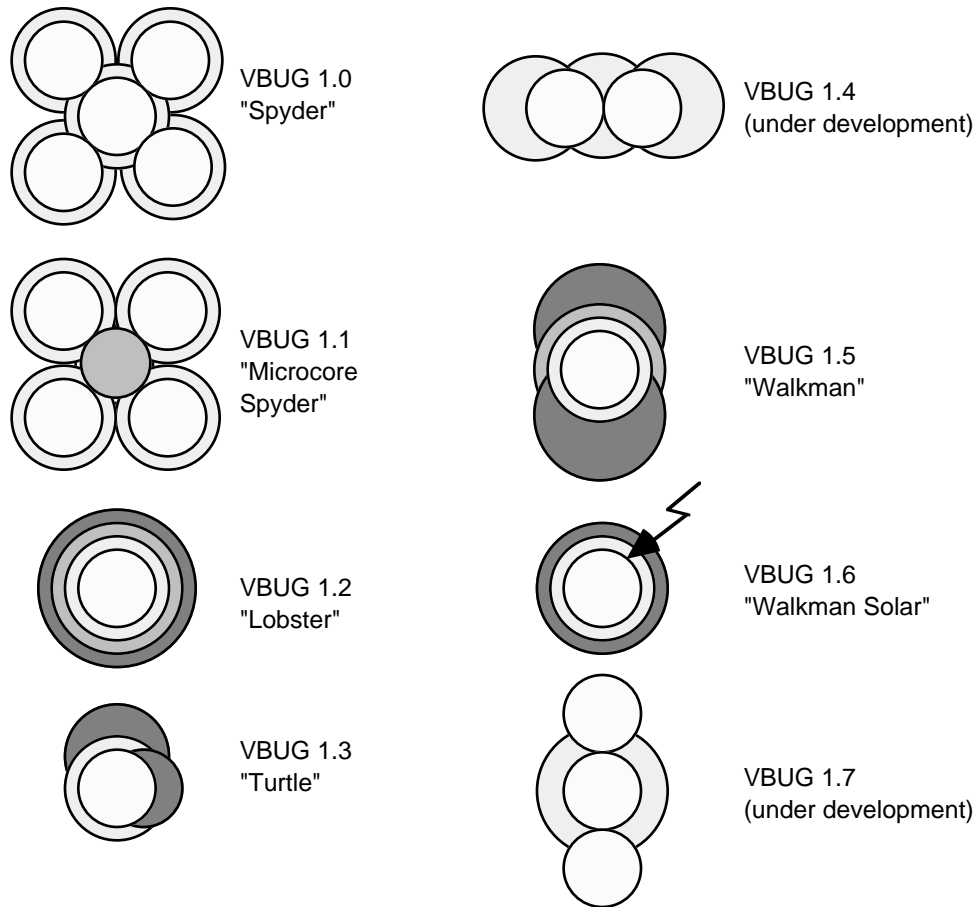
The walkers so far described are the highest biomorphic forms we have so far built, but by no means the only ones. As in biology it is thought that rather than make many phenotypic modifications to a particular generation of device, it is easier to make a device, prove a principle, and when that principle has been phenotypically modified to exhaustion, build another generation. To describe all the resulting designs would be prohibitably long, so Morph diagrams were invented to give a symbolic representation to the variety of designs discussed. Morph diagrams conveniently show the structure of biomorphic layers as well as the Nv mapping which occurs on a particular physical framework. The convention used is detailed in Figure 6:

Figure 6: "Morph" Symbol Diagram Examples



As seen, the morph diagram allows the toroidal biomorphic structure to be mapped onto a plane, allowing for quick and efficient sketching of a wide variety of creatures and their relative abilities. Morph diagrams of the walkers so far built and discussed are detailed in Figure 7.

Figure 7: Morph Diagrams of Prominent Biomorphic Genotypes



Experimental Machine Morphology

From an experimental viewpoint, natural processes have produced such an incredible and capable phase space for life primarily for two reasons: living systems are designed and operate to survive, not to perform blind tasks, and nature is not concerned with the idea of comprehension. An attempt to build machines using these guidelines would appear to be counterproductive, but we submit the opposite is true. By constructing mechanisms capable of immediate and sustained survival we automatically induce in them an appropriate spectrum of behavior to deal with real world situations. They have a basic core survival intellect that can be quickly understood and controlled by classic domestication techniques. Even still, what rules should we formulate as guidelines for designing such survivalist machines?

The working assumptions we use for our machines are twofold:

- For an autonomous machine to have a sufficient survival lifetime one must use resilient materials at small enough scales to have a high weight to power ratio and high structural strength relative to their environment.

The second concerns reproduction and power.

- Machines cannot be made to reproduce themselves easily, nor would we want them to. At small scales, with carbon-based materials platforms, self reproduction could be quite dangerous. Therefore if machines are to survive acceptably then we must extend their lifetime to many years in full operation. This means an autonomous machine must extract power from its environment and since that power may be either weak or scarce, it must be able to operate by storage of and access to power on demand. In normal scale environments this implies solar powered machines. As current solar cell technology is both inefficient and fragile to mount on any mobile design, one must process it through electronic regulation and storage.

From these observations we can extract some experimental rules of behavior for autonomous machines. Unlike logical axioms we consider these malleable rules subject to alteration if experimental evidence implies they are inadequate. First we note that autonomous control has four principal components: sensation, cognition, locomotion and manipulation. For our "primordial" machines, only sensation and locomotion are critical to autonomous machine operation, although it is possible for a species to survive on locomotion alone if it uses hive or herd dependence. We have found the following rules are adequate to ensure the survival of autonomous machines.

Biomorphic Laws

Disregarding other more "esteemed" laws of robotics, the following are rules that will guarantee an autonomous machine's survival.

Biomorphic Survival Laws :

- 1: A machine must protect its existence.
- 2: A machine must acquire more energy than it exerts.
- 3: A machine must exhibit (directed) motion.

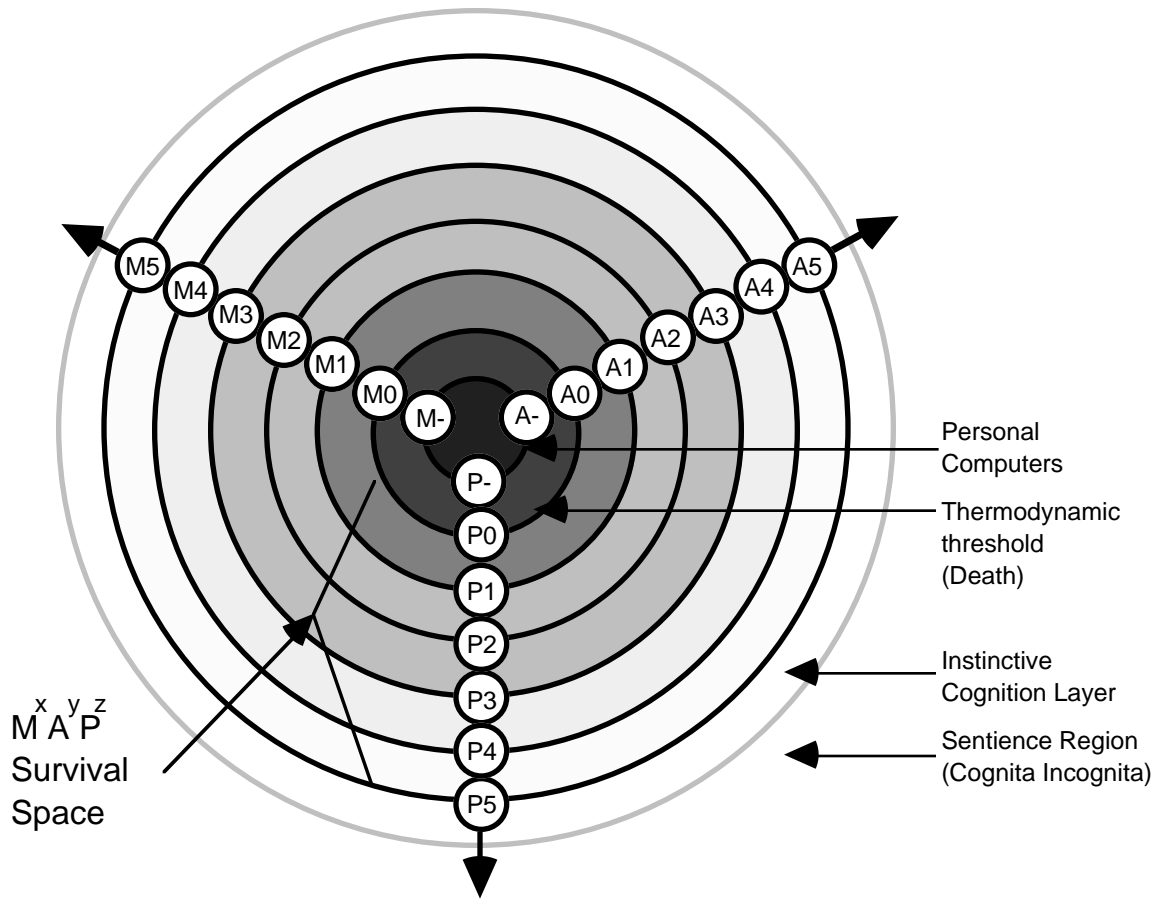
Notice that the survivalist laws are very different from the ethical, and fictitious Asimovian robotic laws. Asimov's laws (in essence, "Protect humans, obey humans, then look after yourself") do make for good fiction, but inadequate survival machines. One or more of the above rules, however, are very easy to incorporate into machine structure and control systems at a minimalist level. Complexity is reduced, and survivability is thus enhanced.

Architectural Maps: "StarNet" Representations

Watching a million dollar autonomous robot bash itself to pieces against a desk edge is a frustrating experience for designers. No matter how much work went into the robot, it failed a basic preservation instinct obvious to any layman. By developing survivor automatons such situations can be averted if not completely avoided, but to do this a survival signature with high resolution must be worked out for biomorphic space. This chapter addresses such a signature.

For the purposes of autonomous biomorphic designs, life is defined as that which moves for its own purposes. This leads directly to the Biomorphic Laws previously mentioned, which in turn form a minimal basis space for a spanning tree of survival capability. Biomorphic Laws can then be encoded in 3 general vectors, Mobility, Acquisition, and Protection (MAP) where each vector length is proportional to a biomorphs capability in that area, and drawn on a planar graph as indicated in Figure 8.

Figure 8: The Complete MAP Survival Space



This diagram places the three Law vectors along the edges of a conical continuous space, increasing from negative survival aspects through layers of exponentially increasing complexity to a high order sentience region. By defining a particular organism's survival proficiency with respect to the Laws, a triangular area can be defined on the space that represents a particular creature's survival metric for a given environment.

Each capability exponent is a milestone of success in a general fractal environment. Consequently many measures will be only part way between states reflecting their degree of survival adequacy. These milestones, taken from experiences in biomorphic studies, are set as:

- M- Motion occurs only under application of an external force.
 - M0 No motion abilities.
 - M1 Moves deliberately in one dimension.
 - M2 Moves deliberately in 2 dimensions.
 - M3 Moves deliberately in 3 dimensions.
 - M4 Capable of dual-mode motion with tools, vehicles, or application specific design elements.
-
- A- Operates from a non-replenishable energy source (battery, power line).
 - A0 Zero energy consumption or delivery.
 - A1 Can directly extract/apply external energy when available.
 - A2 Can efficiently extract/store/utilize external energy.
 - A3 Uses focused tactics to efficiently extract, store, and utilize external energy.
 - A4 Uses planned tactics to efficiently extract/store/utilize external energy.
-
- P- Negative defensive abilities (physically more fragile than environment).
 - P0 zero defensive abilities (structural strength equivalent to environment).
 - P1 flight and/or hide behavior against hostile stimulus.
 - P2 Fight or flight behavior against hostile stimulus.
 - P3 Tactical fight/flight behavior against hostile stimulus.
 - P4 Tool, vehicle, or material use in fight/flight tactics.

Obviously beyond a certain machine capability, survival metrics lose mutual exclusivity. In the animal kingdom, the MAP5 metric (i.e., all survival exponents equal 5) is set as the domain of implicit survival instincts that are observed in all lower animals. Such a metric seems to require advanced, RNA programmed nervous systems that, although not sentient by human standards, gives animals the behavioral tactics to sustain themselves. The flocking of birds, homing ability of whales, turtles, etc., are advanced examples of instinctive, but unconscious, survival strategies. The fact that birds don't avoid airports, and whales fatally beach themselves in their efforts to follow their guidance instinct indicates that these are indeed unconscious, cognitive artifacts.

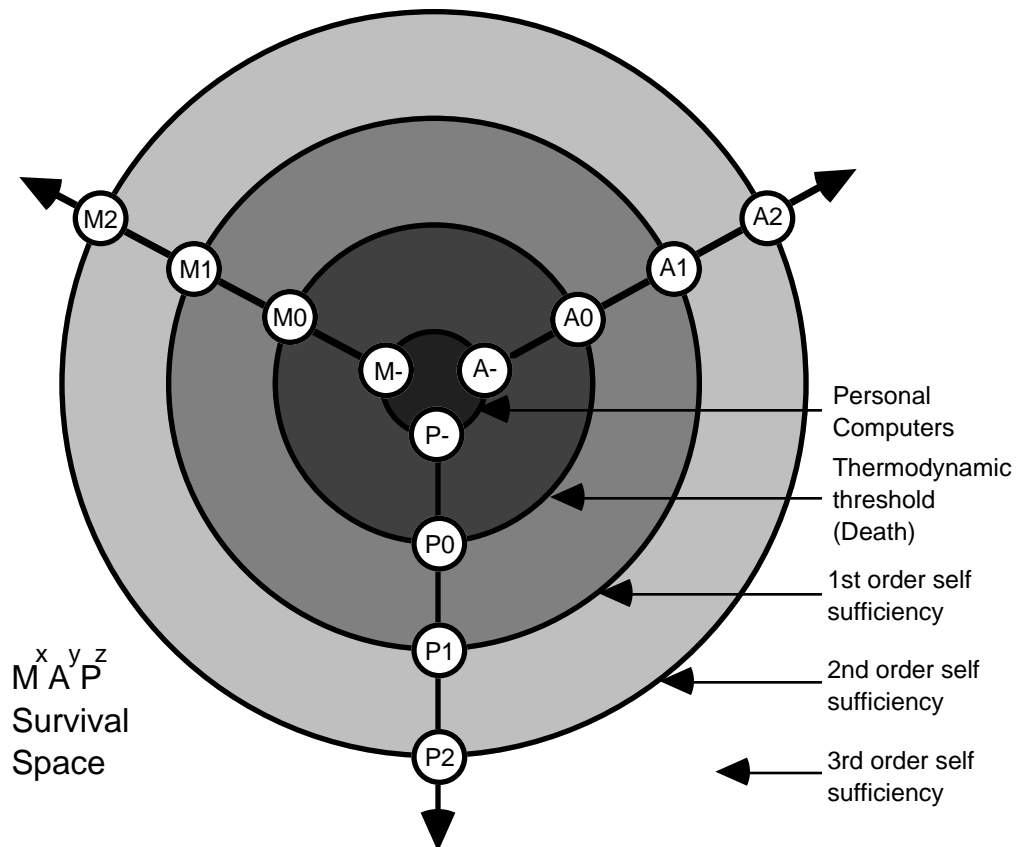
MAP6 is arbitrarily defined as the metric where all survival aspects blend within high level, symbol-based cognitive abilities, and includes the abilities of all large brained animals that can use syllogistic, problem solving logic. Indeed, this realm is where most work has been done in trying to find a functional artificial intelligence (AI). Due to the obvious complexity of this region, we will ignore it until there are advanced enough biomorph mechanisms (and theory) to support high-level AI constructs. We feel that a study of the high-cognitive regions would not be possible until we have a sufficient engineering knowledge of the lower survival dimensions and how to build devices to match them. An indirect aim of this whole technology is to acquire the knowledge of how to build capable mechanistic "bodies". Bodies that can look after themselves so adeptly that any adequate AI construct acting as its brain could concern itself strictly with the problems of world view problem solving, and not, say, how to get its foot out of a gopher hole.

A MAP diagram represents a flexible environment against which biomorphs can be gauged. "Environments" are defined as the dimensional space organisms must exist in, and the consequent metric space for a particular biomorph design will depend on the

application (for example, a mechanical fish would not fare too well on a MAP space for mountain ranges). To be capable at all biomorphs have to be independent, and that requires an ability to exist in general earth environments that are, for all intents, fractal. We assume as the general biomorph environment space the complex surface features found naturally on the Earth's land masses, at scales ranging from 5 cm to 30 cm. This covers biomorph mechanisms that are sufficiently large enough to keep from blowing away due to atmospheric turbulence, but small enough for researchers to avoid excessive design costs. Environment variables must be adequately, if not completely, described for whatever MAP a mechanism must be measured against. For example, at the micron scale nanobots will have to work at, Brownian-motion forces, gravitational fields, strong material densities, even strong and weak molecular forces must be included as crucial environmental variables. Consequently successful "nano-morphs" will not be able to use legs or even wheels for motivation, and their power sources and protective abilities must employ radically different design physics.

So far, biomorph mechanisms have not gone beyond a MAP3 metric, so for the purposes of clarity a truncated MAP space is detailed in Figure 9.

Figure 9: Truncated MAP Core for Simple Organisms

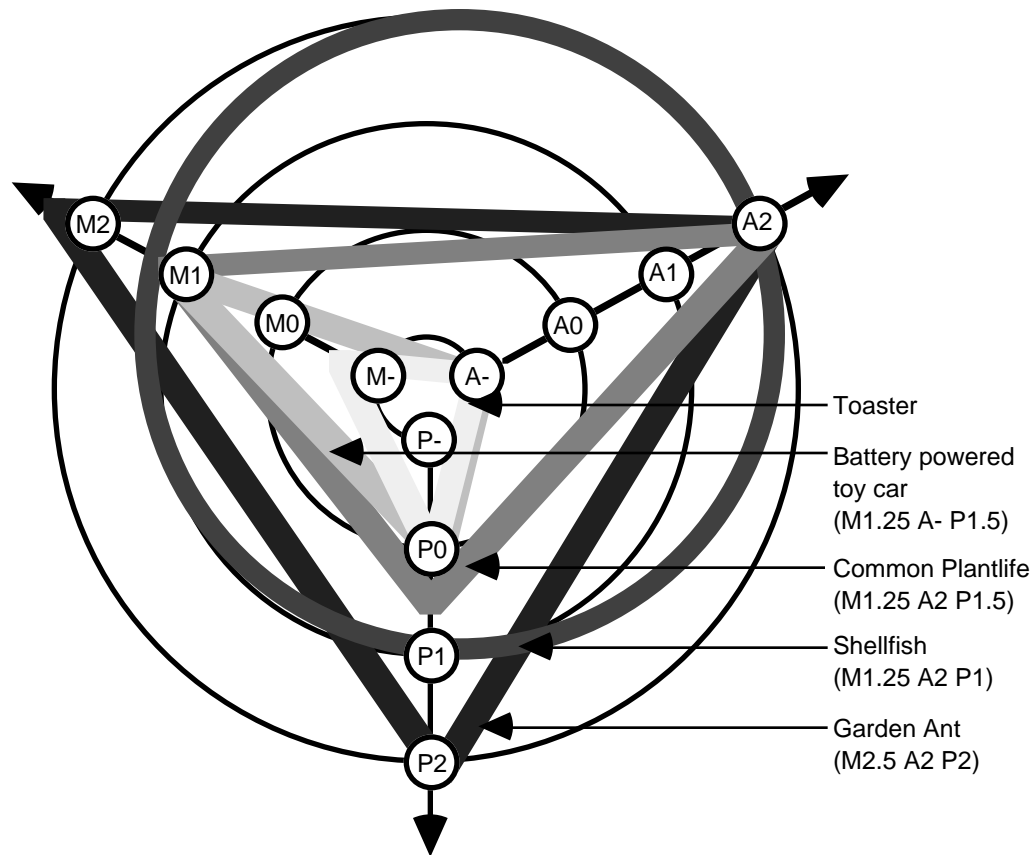


The truncated MAP core shows better detail of the thermodynamic threshold. This threshold marks the boundary at which a device neither moves, feeds, or protects itself, and is the equivalent of death or complete non-function (a uniform MAPO metric would,

for example, represent the survival space of a rock). The metric below this defines negative survival aspects and defines devices that must be moved, fed, and protected from outside sources. A good biological example of a uniform MAP- (MAP-negative) metric would be a bird's egg, and a good technical example of would be a personal computer.

As suspected, the MAP represents a survival space upon which can be measured not just biomorphs but most simple biological organisms, machines, and even children's toys. As such the MAP is useful to see how biomorphic mechanisms rate against these other biomimetic devices. A MAP diagram showing some common survival vectors is detailed in Figure 10:

Figure 10: MAP Vectors for some Common "Life" Forms



The vertex of each triangle converges at the point where a particular creature's survival scale is measured. As the three main MAP vectors abound in fuzzy logic connections, for many creatures there are half and quarter way points along each. For example, most plants are harder than the immediate objects in their vicinity, but do not use flight-or-hide behavior against aggression. This would give them a P0 rating if it were not for the chemical and/or thorn defenses most plants employ, which raises them to a P1.5 protective rating. Such labels are chosen within a particular MAP space relative to other creatures measured, and so the metrics are broad values that are somewhat subjective until a sufficient database has been established. As the resolution on biomorph-like metrics can be made fairly fine however (i.e., ants obviously have a

much larger survival spectrum as compared to plants at regular time scales) we believe that such a labeling scheme will be more than sufficient for comparison purposes.

Survival Signatures

The volume inside a survival triangle could be calculated to generate a single survival vector but this would be a poor scale for two reasons; one, It is possible that two creatures could have the same scale value though they might have radically different survival spectra; two, some creatures (like the shellfish) have a bizarre metric where all survival vectors are shared within a single design feature. As shellfish use their shells for protection, feeding, and (in some species) propulsion, any triangle placed upon the MAP space can be rotated uniformly between survival axes, resulting in the circular survival metric shown. Shellfish MAP space is thus indistinct, but shows how MAP diagrams are useful for detailing a broad range of genotypes.

A survival signature can be calculated, however, if we assume higher hierarchy of weights for each of the MAP scales roughly proportional to their importance. To keep the resulting signature linear (despite their exponential implementation complexity), we assign M=3, A=2, and P=1. The equation is:

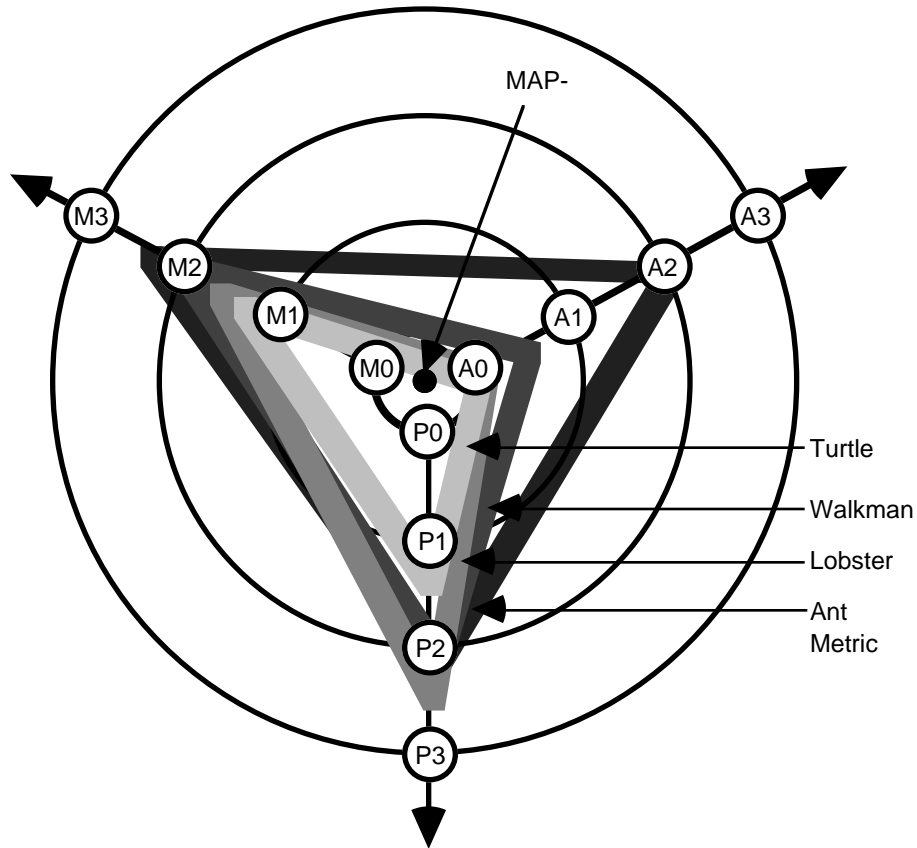
$$\text{Survival Signature Space (SSS)} = 3 \times M. \text{ metric} + 2 \times A.\text{metric} + P. \text{ metric}$$

So a survival signature for an ant would be 13.5. For creatures with distinct MAP metrics, this signature is fairly accurate. For more indistinct creatures like the shellfish, further study is necessary. Fortunately as this system has been devised to measure our pseudo-linear biomorphic devices, this measure has shown itself to be quite adequate.

By convention, when MAP values are between states they are represented by a multiple of quarter vector lengths. This provides sufficient resolution to describe a creatures' ability without being absurd. An example would be the motion vector (Mx) for a toaster. If the toaster is a standard one, we apply a motion force that it eventually returns by release of a spring, thus rating it at M-0.5. If we had a "soft-touch" toaster that used a motor triggered by pressing an electronic button, we could assign it an M-0.25 rating, as the mechanism is using some sort of "smarts" to obey our command. This is obviously higher than the standard mechanical model, but below the M0 rating, as we had to begin the action externally. If however, the toaster is defective, so application of force yields little return (i.e., stuck toast), the motion metric drops to M-0.75. This leaves only the M-1.0 or M- metric to describe toasters where no matter how energetically you pump the lever, you'll be eating cold cereal that morning.

As we generally assume biomorphs to have a greater survival signature than a toaster (SSS = -3.5), the detail region within the energy threshold can be shrunk to a point. This gives a reasonable resolution to the third survival exponent and increases the resolution for any displayed metrics. An example MAP space, showing the survival spaces for the walkers previously described as compared to a common garden ant is shown in Figure 11. The Spyder metric is not shown as its indistinct space has yet to be classified.

Figure 11: MAP Space for 3 Biomorph Walkers as compared to a Garden Ant



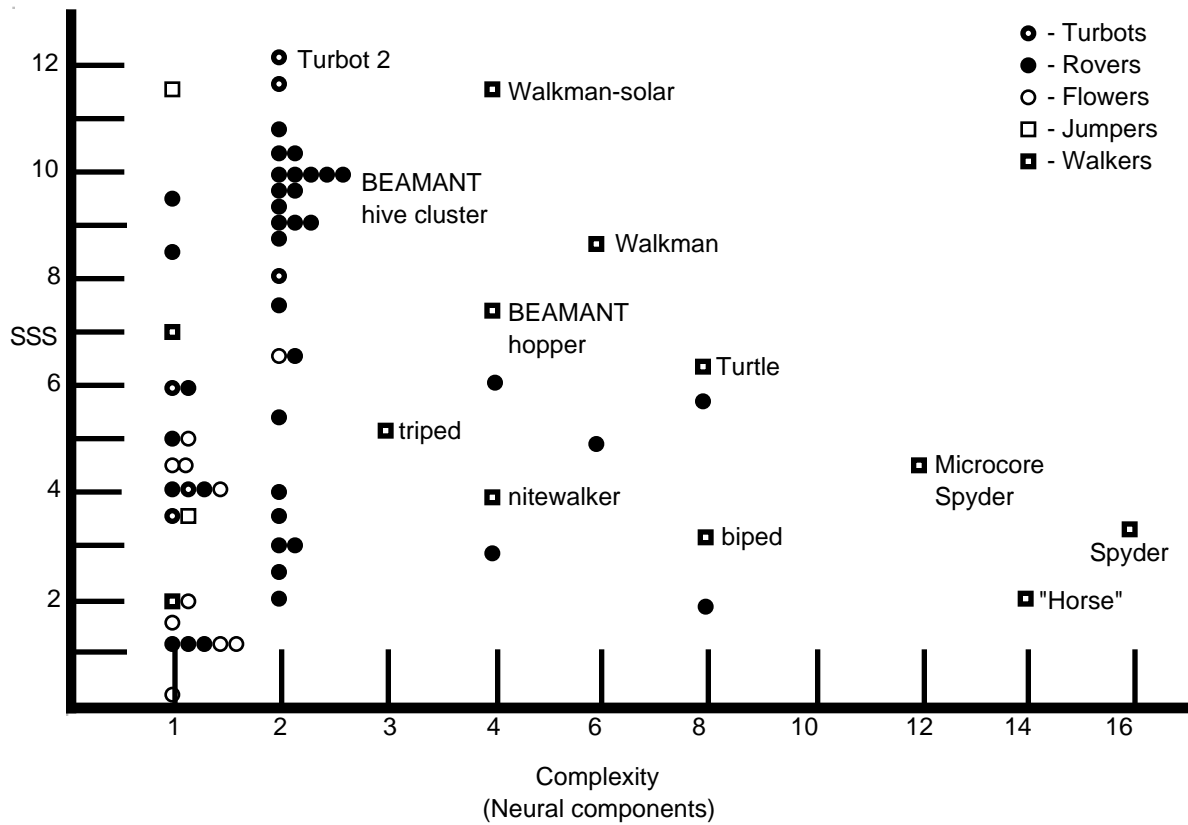
The common garden ant metric (M2.5 A2 P2.5) is a measure against which we compare biomorph designs because it is commonly recognized, easily understandable and is a design ideal our mechanisms aspire to, at least for the present. As an explanation of the ant metric, M2.5 refers to the ant's ability to handle 2 dimensional travel with ease, though ants can handle three dimensional terrains just short of jumping (2.75), which is just short of flying (3.0). A2 refers to the individual ant's ability to ingest, process, and store hive sugars for fuel. It does not refer to the ant's participation in the hives collective ability to process foraged food. Many ants do not eat the foodstuffs they find, only return it to the colony where it is converted into manna that the ants ingest easily. As such, individual ants only eat and store their food, they do not "hunt" the manna anymore than humans hunt a loaf of bread, so they only rate a second order acquisition metric. The protection metric P2 refers to the ants simple tactics of either flight from an unknown enemy, or fight by charging opponents it is chemically adverse to.

It is interesting to note that an ant colony (considered as a single organism) would have a larger survival space for it gains from the emergent properties of massed individual efforts, but that will have to wait for another paper. Right now, we have a tool and examples against which we can justify the survival success for the 70 or so mechanical biomorph designs currently under study. The designs discussed so far rate as follows:

"Turtle" M2.25 A- P1.5, SSS = 6.25
 "Lobster" M2.25 A- P2.25, SSS = 7
 "Walkman" M2.5 A-0.5 P2.0, SSS = 8.25
 "Spyder" (approximately 4.5)

But taking the many devices so far devised (some of which are shown in the next section) we can make a complexity/survival graph as follows, which shows the evolutionary progress so far attained.

Figure 12: Survival Signatures verses Complexity for 57 Active Biomorph Devices



Complete biomorphic device details are not included here because many devices have not yet undergone their final phenotype modifications, but also because this is a paper, not a book. Most solar-powered biomorph devices will improve based upon the rate of new, applicable techniques, and their comparable survival abilities. As can be seen, microcore designs improve survival characteristics, peaking at Walkman-Solar with almost an equivalent Turbot survival space. Full details will be included in future papers.

Neural Morphology

For biomorphs past a certain level of complexity, intelligence emerges as a collective effect by interacting with complex environments. We now look at design constraints and advantages to optimize our designs.

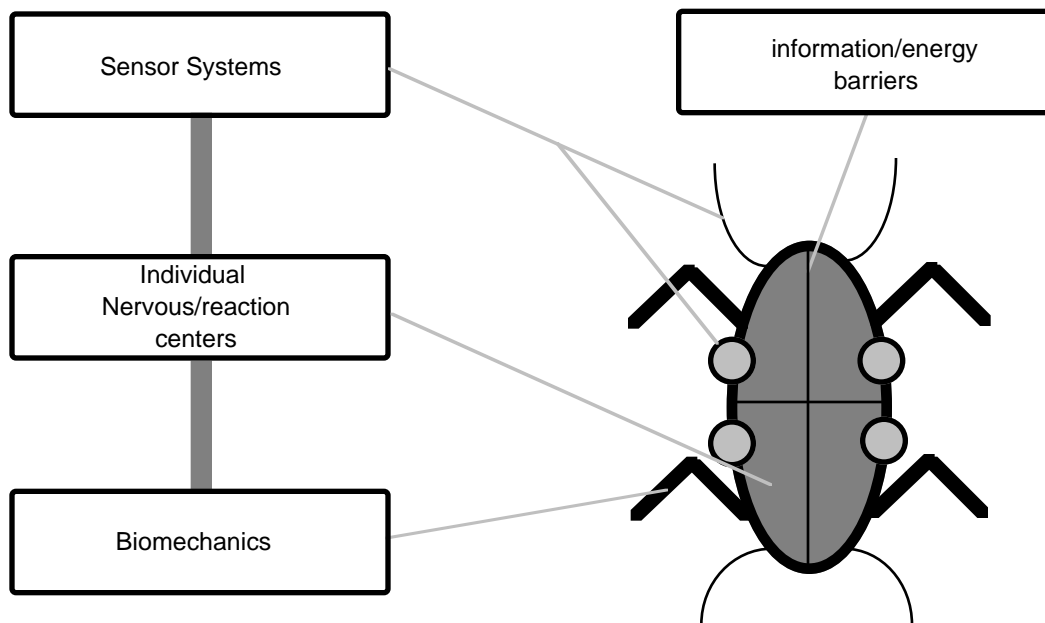
The most successful biologic survival tactic, breeding, cannot be used for machines because of their chemical makeup and the incredible energies necessary for the task. Fortunately for this argument, we assume a human as a machine's way of making another machine. The qualifier for this is that anyone who makes copies of the same machine is just a reproductive mechanism, whereas anyone who builds a new machine as an improvement on a previous machine "genotype" can be considered a force of directed evolution for that species.

When a machine breaks it effectively dies, so another aspect of the CA design structure is to build robots that have a significantly long life span. This is useful as survivalist elements can be observed and studied over a long time scale, and against newer generations of biomorphic life.

There are three general classes of biomorph design; invertebrates, vertebrates, and cooperative organisms.

Invertebrates represent clusters of quasi-independent mechanisms within a single mobile chassis. Relationships between these sub-elements range from tenuous to direct electrical and mechanical linkages, resulting in low element count but high survival indexes. A modular diagram is detailed below:

Figure 13: Invertebrate Control Structure

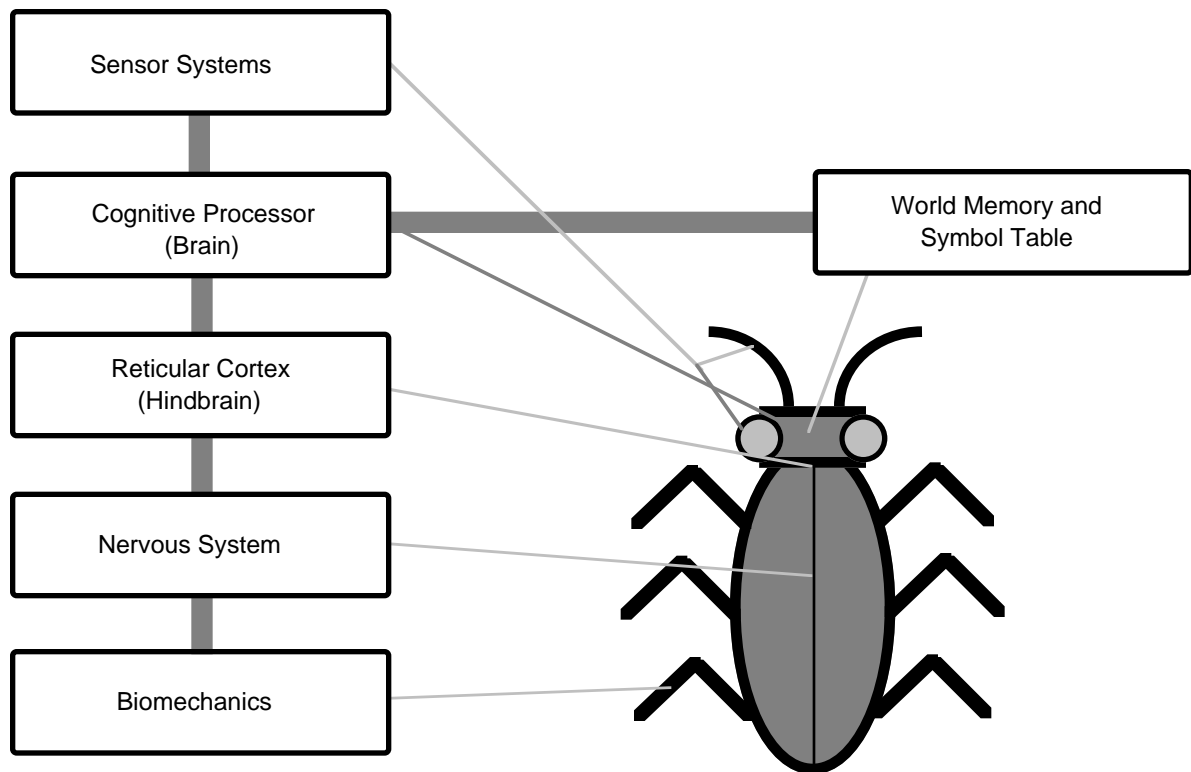


These are the least developed of the biomorph designs so far because they are the most counterintuitive. They will be the subject of future papers as soon as a wider machine

spectrum has been built and studied. Though the continuously running Robot Jurassic Park (where most devices built so far are continuously interacting) has had a few surprises, none was so obvious as the 4 transistor "Turbot 2" (M2.75 A3 P2, SSS = 16.25) which exhibits aggressive phototropism to the point where it will systematically try every way to get over a significant obstacle to a brighter light environment. It has no focusing apparatus, yet in a general lighted environment it exhibits strong phototaxis. It is a two-neuron creature in a capable, point symmetric body, and is so far the undisputed "robotus-rex" amongst its weaker cousins. Turbots are the subject of ongoing research, and amazingly capable for an invertebrate design.

Vertebrates are structured around a concentric spine through which sensor, power and other information flows, usually in a top down hierarchy, and in a bilateral arrangement. Vertebrates use their topological advantages to synchronize actions between actuators (i.e., drive motors) and sometimes sensors. All biomorph walkers, with the exception of microcore Spyder, work from subsets of this arrangement.

Figure 14: Vertebrate Control Structure



In biomorphic walkers, the most successful arrangement has been to design from the mechanics back. As nature has proven time and time again, good controllers never make up for inadequate mechanics, but knowing at least the network properties that will control our creatures allows us to make broad assumptions about the mechanics we employ. Biomorph design, because of the flexibility of the controller, allows for asymmetrical structures that most synchronized controllers would not tolerate. This is a major advantage in device construction as designers can build devices with vastly different leg styles, balance centers, and suspension on the same chassis and still expect

to get an efficient convergent solution. Also, asymmetrical designs are much more interesting, easier to build, and there is experimental evidence they can be inherently self-stabilizing.

Vertebrates are interesting from a biomorphic view because not all morphic layers are necessary to make a sufficient design. Mechanics and a nervous system have been shown, by experimental evidence, to be enough. Anything beyond this basic core enhances the design's survival space, but is optional baggage. This is seen in biological examples in that most life forms survive quite well with a lot fewer neurons than there are transistors in a pocket radio. The implication is not that designers can substitute smaller and cheaper controllers, but that survival skills can be based upon much simpler precepts than world-model symbolism. After all, the first primordial creatures couldn't possibly have been a brain that evolved a body, but the reverse, and remains so for the majority of all known living organisms today.

We conclude that making biomorphic machines conform to anthropomorphic ideals is not a good idea for cases where survivability is essential. Robobiology can be based on some biologic examples (i.e., insects), but this is because insectoid life is probably the most mechanical of all life forms, not the other way around (an example of anti-anthropomorphology).

Experimental Methodology

The final biomorph form, cooperative organisms, is the most elusive of the designs to date because it requires a large array of diverse machines and the space to allow them to interact. So far there are only 60 or so working agents ranging from rovers, walkers, jumpers, spinners, and tumblers under study on a table the size of a standard office door. Hives exhibit greater abilities than individual elements, which is well known, but whether this can be extended to biomorphic mechanisms is not clear. Biomorphs do not need to socialize for basic survival and so there has been little evolutionary force to make them do so. In the Robot Jurassic Park (RJP), where over 40 robots of 12 different solar-powered species have been running continuously for over 6 months (as of the time of this writing), there has been evidence of flocking, fighting, cooperative group battles against particularly aggressive forms, even pecking-order dominance, but little in the way of true cooperation that would indicate hive structure stability for such devices. It is suspected that further work will have to be done in sensor technology so that like creatures would be able to recognize others of their own hive.

A major lesson from the park is that among different, selfish species, cooperation is not only possible, it is inevitable. This is seen repeatedly as the simple, two-to-four neuron creatures in the park exhibit a vast array of recognizable biomimetic behaviors, usually as a result of instigation from machines of a completely different design.

Figure 15: A Moment in the Life of the Robot Jurassic Park



This picture shows 15 biomorphs in interaction, composed of BEAMANT rovers, "TURBOT" variants, solar walkers, and a three legged "hopper", none of which use more than two Nv neurons. The two-Nv design is the minimal necessary to make a planar rover, where each neuron powers a single motor in a closed chassis. The range of 2 Nv creatures developed (over 30 species, one of which is the vicious Turbot 2) shows the principle of biocognitive intelligence. That is, survival traits are determined by the physical structure of a design. A classic example is evolving offset visual sensors so that phototropic photovores (light-seeking light-eaters) will not dive straight into a damaging fire, but circle it like a moth.

So our design philosophy is to make a variety of self-contained minimalist robots, building on advantages in design observed from each generation of machine. First, we get them to survive, later, we train them to do tasks. Right now, work is concentrating on introducing "tilebot" social designs, a whole new species genotype, into the park. How this will affect the current loner devices in constant interaction should be interesting.

Implications of Scaling and Modularity

Biomorphic architectures are not confined to ordinary scales, and fascinating things happen when we consider extraordinary scales for these machines. We now focus on downscaling to micron scales and below, to the nanoscale. At small scales, biomorphic structures would not look like ordinary bioforms, but all the principles of biomorphic architectures can be fully used.

Downscaling requires two critical properties: The first is architectural scale invariance, especially in the use of analog computation; The second is system design modularity; embedding the organism in StarNet fashion physically as well as abstractly. Without scale invariance and modularity, extending these principles to other scales would require developing new concepts, but with the biomorphic architecture previously described the transition is smooth.

Micron Machines

The first class of thought machines, constructable with current technology, is descending to a few tens of square microns. Once we see how to build a simple but effective biomorphic machine at this scale, we can discuss a further shrinking of machines to the order of a square micron or slightly below. This is the size of a typical cell, a few microns in diameter. At this scale self-assembling colonies of biomorphic machines are not only practical, but desirable. Colonies should both self-organize physically and develop collective or emergent behavior (as the loner or social creatures under study do now), which means that the collective machine has properties that each modular piece does not have. These meta-scale ensembles are complete organisms; super-machines that have many interesting properties and uses.

The powerful tool of silicon chip technology is available at micron scales, and we will use it to essentially print machines. Lithography can etch machines in parallel quite cheaply. The simplest micron scale medium is a liquid that we take to be water. Instead of a walking machine, we will design a swimming machine, again using many clues on how to do this effectively from biological organisms (indeed, at such scales the viscosity of the environment would be so high that legs, transformed into cilia, may be the only practical means of directed locomotion). The design must be both minimal and respect the constraints imposed by silicon fabrication on materials and geometric form.

Using autonomous mobile machines of biomorphic architecture, we construct micron size colonies of autonomous micromachines and examine their self assembling properties. Equipped with simple oscillator-driven, mechanical micromachined "flagella" (operating in a linear mode) and a light driven power supply, these micromachines can be used (i.e., tricked) to perform a wide variety of tasks.

A natural mechanical drive is an oscillating flap or lever, micromachined with MEMS technology. A hybrid assembly of the drive is attached to a photocell with simple circuitry to control a charge/discharge cycle. Driving power is not continuous, but uses the demonstrated strategy of larger models of "store until able to move". The biomorph would negotiate a fluid environment. A practical initial size is roughly 100 microns square. Initially the hybrids will have to be hand assembled using STM tools, though eventually they can be printed. We can then study the collective behavior of a colony of real biomorphic machines as a first step toward biological hybrid approaches (as well as adequate simulation). The advantages are threefold; the colonies can be vast in number, the exploration area unlimited, and the size of the devices would mean that interactions could occur at far greater rates than are possible at conventional scales.

Now split the micro machine problem into two parts, a part concerned with dynamic parallel self assembly of small tiles and a part concerned with propulsion and motivation of individual tiles.

Using the Cell

The major problems currently blocking the construction of nanorobotic devices are how to make controlling computers that will fit in a cubic angstrom, and how to get the power the devices will need to work. We suggest that biomorphic architectures may solve the

first, and that using the respectable potential difference across biological cell walls will provide the second.

Micron Machine Colonies - Super Machines

To self assemble machine colonies we first make sub micron sized silicon tiles with five and six sides (three sided macroscale prototypes have been built and are under study now). We want them to self assemble on a roughly two-dimensional surface with three constraints: that the patterns that emerge be controllable and varied; that there be special tiles that can be used for central control; that there be enough power available to allow us to use simple thrust engines etched on each tile. For this it helps to use the wide availability of biological cells to aid assembly.

The model we use is a dynamic soccer ball covering which takes advantage of intrinsic geometry, topology and the organic functional groups already developed in the nanoscience community for self assembling wires. What we are after is not a single machine (though we must start there), because at the micron scale the biological environment is too harsh for a (single cell)/(single function) machine organism to survive. Colonies that dynamically glue into more complex machines offer greater functionality and survivability.

Consider an isolated, roughly spherical cell. Using standard Euler arguments we know that a topological sphere requires 12 pentagons and the rest hexagons to tile it completely. We take the pentagons as special, identifiable by their five sides and the hexagons as indistinguishable slave tiles that can provide propulsion. Now we use the same tools that self assembling wires use -- organic glue, and selective functional groups to do the self assembly. The replacement for the target metal or bandgap pads normally used in self assembling nanowires is the cell surface itself. Selectivity of chemical groups is focused on the edges and surfaces of the material used which we will take to be silicon (though a plastic would serve just as well). One uses batch chemistry to attach a selective layer to one side of a tile that attaches to the surface of the cell, but allows the tile to slide over the cell's surface.

Functional groups are put on the edges of the tiles which one can imagine have several colors, say red-green-blue-etc. Attraction occurs for like colors (red-red) and repulsion for different colors (red-green) and various permutations of this scheme. This way we avoid immediate lock-up of the tiling and allow it to take on various configurations depending on what functionality we wish it to have. Further sorting aspects emerge as a function of the edge geometry of the tile edges, as well as functions of dedicated cilia that allow the tiles to flexibly "velcro" together once alignment has been obtained. This again can be done with functional organic groups.

This is self assembling tiling done in parallel and will work provided we do not kill the cell with toxic functional groups. This will be somewhat deferred, however, as the tiles themselves will each have their own marginal (possibly solar) power source that allows them sufficient autonomy to survive ($SSS = 11.0$ assuming a 3d motility environment), and directionally self-tiling. Note that the range of functional groups available is much wider than that for the self assembling wires currently under development since we do not require conduction, but simple polar bond adhesion. A problem for which there are many known biochemical solutions.

To get a machine colony instead of a tiling complex, we enable the tiles by making them simple biomorphic machines suspended in solution. Conversion of a tile to a mosaic machine needs three things, a power source, a control mechanism, and propulsion. The architectural topology for the tile would most naturally be an extended Turbot topology of biomorphic architecture, which is basically two simple chaotic oscillators driving separate flagella of unequal lengths, weakly coupled by a two-transistor neuron "brain" circuit. It is the simplest multi-cellular machine organism with capable, complex behavior and can be likened to a machine virus. Once a cell is tiled, power now comes not from light, but from the considerable potential difference across the cell membrane (approximately 400mV). This can be extracted with pronged electrodes that pierce and adhere to the cell membrane, or a redundant flagellum on our micromachine designed for the task.

Biomorphic machine colonies are suspected to exhibit a wide variety of complicated emergent behaviors from simple seek and avoid, to cooperative tasks such as foreign body rejection, to super cluster colony construction. This scenario is a case study -- there is nothing present in this scenario that we cannot do now if we wished, modulo the proper power extraction mechanism for the cell, which, though forefront, is in the realm of biologists. The active tiles or mosaic machines could be made very cheaply by standard techniques, since cost to fabricate goes up with continuous area and the area of these discrete tiles is very small.

There are several features of these self assembling active mosaic machines that are crucial: the assembly is massively parallel, and the whole concept scales downward to the nanoscale in a clean way. We can simulate these devices with computer models, and build large scale versions in a simple water tank as a proof of concept. This work is under way.

It is interesting to note that a colony of diverse biomorphic tiles could create a space of finite elements that organize into larger collective creatures with a potential for tile-cluster "reproduction". That is, there may be tile structures that make copies of themselves from the suspended tile matrix they "live" in. This is purely speculative but it would constitute a dynamic proof of the Von Neumann self-replication principle, and would lead to observed machine reproduction in a safe, linear regime, rather than dangerous exponential growth. If the technology can be developed, it is possible that such "breeders" might emerge from a sufficient biomorphic tile space. As with similar genetic-algorithm projects, we can let the devices emerge on their own, or we can deliberately design them once we understand their behavioral characteristics (as with "game of life" constructs). Though such a breeding scenario sounds ideal for computer simulation, as this paper has hopefully shown, results could come probably much faster and more effectively if these tiles existed in reality.

Nanoscale Machines - Inside the Cell

Generalizing the example above we see that we can do without a cell. The cell was a convenience put there for both an assembling surface and power; these machines will form self organizing colonies anyway. As we descend to 100 nanometers and below we can go inside the cell, which has several advantages. There is abundant ATP in the cell and many styles of available ATP engines that seek ATP gradients and use them for locomotion. If we can find functional groups that would selectively adhere ATP engines to a small substrate, we solve at least the power problem. Control is more subtle since circuits become difficult to build as we descend in size, but a two-transistor circuit

should be feasible where larger designs would not. We also need mechanical flagella, which is not a problem, but an oscillator driver might be hard. There may be a much more clever way to gain functionality inside the cell using biological mechanisms. If we can do this scale of reduction and reach the inside of the cell we have a very powerful tool; i.e., a fixed amount of inert, mechanical force that can be used as either an active catalyst or suppresser. We use the cell to protect the colony; move freely within the cell; operate on or modify internal cell structure; gain free power from glycolytically driven ATP engines. Self assembly at this scale is already massive: ATP engines self assemble by biology; empowered machines can self assemble into mobile colonies with mechanical flagella and complex but targetable behavior, depending on the coupling between the machine halves.

Molecular Self-Assembly

The cell-nanoscale picture would be complete if we could learn to self assemble all machine components, including a simple neuron system. Flagella can be made from self-assembled organic beams which can be made very rigid. We have organic staples at our command from wire research so we can auto-glue almost any required mechanical configuration. Simple oscillator pancake motors can be made self assembling as can capacitors for power storage. It would be even more elegant if we borrow self-assembled nanomotors used by bacteria to drive flagella and an ATP based power source for them. All we need then is the ability to make a simple self assembling neuron control system. As we learn more about building nanomechanical structures and self-assembling wires and components, options will become obvious that we have no access to now. That is a way off, but in a sense working at our current design scale is the key to a huge variety of applications that are quite natural once we come to understand such structures.

We summarize some key points about such micro and nano machine colonies that will be useful in other contexts. We began by a simple extension of current research on the self assembly of organic wires and looked at the feasibility of machine colony self-assembly, descending downwards in scale. The structure of individual machines is elementary. We use the concept of organic glue with attractive and repulsive interactions to do dynamic self assembly on a surface. If the surface is a membrane, we try to extract power from it. These machines could just as easily get power in other contexts inductively or from direct contact with a power or signal bus. These machines are autonomous which means they require no instructions from the outside world, but their behavior is, if not predictable, bounded so that we can analyze and control them effectively. They are mobile and their drivers can be very simple oscillators drawing little power to move mechanical flagella, which are also simple. A Siamese-twin loose coupled architecture is very rich but there are many other possible biomorphic architectures as yet unexplored. Control circuits as synthetic neurons, as we have shown, can be very minimal. Colonies of such machines exhibit very complex behavior. Solutions have presented themselves as minimal, elegant, and accessible. It requires only a paradigm shift of the "robot" concept as depicted, and researchers willing to take up the challenge.

BH
MWT
June 14, 1994

Bibliography

- Anderson, B.D.O: "Stability of Control Systems with Multiple Nonlinearities", J. Franklin Inst., Volume 282, Number 3, 1966.
- Anderson, T., Donath, M., "Animal Behavior as a Paradigm for Developing Robot Autonomy", Elsevier Science Publishers B.V (North-Holland), P 145-168.
- Arkin, R. 1990, "Integrating Behavioral, Perceptual, and World Knowledge in Reactive Navigation", North-Holland Robotics and Autonomous Systems, Vol 6, P 105-122.
- Arkin, R, "Intelligent Mobile Robots in the Workplace: Leaving the Guide Behind", ACM Press, The First International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems, IEA/AIE-88, Vol 1,1-3 June 1988, P 553-561.
- Arkin, R, "Integrating Behavioral, Perceptual, and World Knowledge", North-Holland, Robotics and Autonomous Systems, Vol 6, 1990, P105-122.
- Arkin, R, "Motor Schema Based Navigation for a Mobile Robot: An Approach to Programming by Behavior", IEEE Proceedings of the IEEE International Conference on Robotics and Automation: REPRINT, 31 March - 3 April, 1987, P 264-271.
- Atherton, D.P.: "Stability of Nonlinear Systems", Chapter 3, Research Studies Press, England, 1981.
- Babloyantz, A. and Destexhe, A. (1987)"Chaos in Neural Networks". In: Caudill, M. Butler, C. (ed.) IEEE Intl. Conf. on Neural Networks (1st) 4:31-39.
- Basti, G. and Perrone, A. (1989) "On the Cognitive Function of Deterministic Chaos in Neural Networks". IJCNN Intl. Joint Conf. on Neural Networks 1:657-663.
- Bitsoris, G. et. al.: "Stability Analysis of Complex Discrete Systems with Locally and Globally Stable Subsystems", Int. J. Syst. Sci., Volume 25, pp. 911-928, 1977.
- Blum, E. K. and Wang, X. (1992) "Stability of Fixed Points and Periodic Orbits and Bifurcations in Analog Neural Networks". Neural Networks 5: 577-587.
- Braitenberg. V., "Vehicles: Experiments in Synthetic Psychology". MIT Press, Bradford Books, Cambridge, MA, 1984.
- Brockett, R.W.: "Frequency Domain Stability Criteria", IEEE Trans. Automatic Control, AC-10, pp. 255-261 and 407-413.
- Brooks, R, Connell, J, Flynn, A, "A Mobile Robot with Onboard Parallel Processor and Large Workspace Arm", The 5th National Conference on Artificial Intelligence, AAAI-86, V 2, 11-15 August 1986, P 1096-1100.
- Brooks, R, "A Robot that Walks; Emergent Behaviors from a Carefully Evolved Network", IEEE Computer Society Press, Proceedings of the 1989 IEEE International Conference on Robotics and Automation, Vol 2, 14-19 May 1989, P 692-694.

- Brooks, R., "A Robust Layered Control System for a Mobile Robot", Massachusetts Institute of Technology, Artificial Intelligence Laboratory, MITAIL, A.I. Memo 864, September 1985.
- Brooks, R, "Engineering Approach to Building Complete, Intelligent Beings", SPIE - The International Society for Optical Engineering, Intelligent Robots and Computer Vision, Editor: Casasent, D, Vol 1002, 7-11 November, 1988, P 618-625.
- Brooks, R, "The Whole Iguana", MIT Press Robotics Science, Editor: Brady, M, 1989, P 432-456.
- Chapeau-Blondeau, F. (1993) "Analysis of Neural Networks with Chaotic Dynamics". *Chaos, Solitons & Fractals* 3:133-139.
- Chapeau-Blondeau, F. and Chauvet G. (1992) "Stable, Oscillatory, and Chaotic Regimes in the Dynamics of Small Neural Networks With Delay". *Neural Networks* 5: 735-743.
- Chen, B-R., M.J. Hines, and H. Hemami. "Dynamic modeling for implementation of a right turn bipedal walking". *Journal of Biomechanics*, 13(3):195--206, 1986.
- Cohen, M. A. (1991) "The Construction of Arbitrary Stable Dynamics in Nonlinear Neural Networks". *Neural Networks* 5:83-103.
- Connell, J, "Creature Design with the Subsumption Architecture", Morgan Kaufmann, Proceedings of the Tenth International Joint Conference on Artificial Intelligence, IJCAI-87, Editor: McDermott, J, Vol 2, 23-28 August 1987, P 1124-1126.
- de Garis, H., "ECAL93 REPORT", *Alife Digest*, Number 107, Wednesday, July 7th 1993, P 13-19.
- Dennett, D., "Consciousness Explained", MIT Press., 1992.
- "Emergent Computation" (edited by Stephanie Forrest), The MIT Press, ISBN 0-262-56057-7 reprint of a special issue of *Physica D*, Vol 42, 1990.
- Freeman, W. J. (1990a) "Nonlinear Neural Dynamics in Olfaction as a Model for Cognition. In: *Chaos in Brain Function*". Baar, E. (ed.) Berlin: Springer-Verlag, 63-73.
- Hodgins, J. K., "Biped gait transitions," 1991 IEEE International Conference on Robotics and Automation.
- Hogan, N., "The mechanics of multi-joint posture and movement control". *Biological Cybernetics*, 52:315--331, 1985.
- Holden, A. V. and Fan, Y. (1992) "From Simple to Complex Oscillatory Behavior via Intermittent Chaos in the Rose-Hindmarsh Model for Neuronal Activity". *Chaos, Solitons & Fractals* 2: 349-369.
- Holtzman, J.M.: "Nonlinear System Theory - a Functional Analysis Approach", Prentice-Hall, Eaglewood Cliffs, New Jersey, 1970.
- Jones. J.L., Flynn. A.M., "Mobile Robots: Inspirations to Implementation", A.K. Peters Ltd., ISBN 1-56881-011-3, 1993.

Kadonoff, L, Benayad-Cherif, F, Franklin, A, Maddox, J, Muller, L, Moravec, H, "Arbitration of Multiple Control Strategies for Mobile Robots", SPIE Mobile Robots, Vol 727, Oct 1986 (published 1987), P 90-98.

Kaelbling, L, "An Architecture for Intelligent Reactive Systems", Morgan Kaufmann, Proceedings of the 1986 Workshop: Reasoning about Actions and Plans, Editors: Georgeff, M, Lansky, A, 30 June - 2 July, 1986, P 395-410.

Khalil, H.K.: "Nonlinear Systems", MacMillan Pub. Co., New York, 1992.

Kowalski, J. M., Albert, G. L., Rhoades, B. K. and Gross, G. W. (1992) "Neuronal Networks With Spontaneous, Correlated Bursting Activity: Theory and Simulations". Neural Networks 5: 805-822.

Letov, A.M.: "Stability in Nonlinear Control Systems", Princeton University Press, New Jersey, 1961.

Maes, P., "Building Artificial Creatures", ARS Electronica 93, Proceedings, LINZ Press. 1993, P 184-197.

Mahmoud, M.S. et. al.: "Large-Scale Control Systems: Theories and Techniques", Chapter 4, Marcel Dekker Press, New York, 1985.

McFarland. D., Bosser. T., "Intelligent Behavior in Animals and Robots". MIT Press, ISBN, 0-262-13293-1, 1994.

McGeer, T., "Passive dynamic walking". The International Journal of Robotics Research, 9(2):62--82, 1990.

McGeer, T., "Passive walking with knees," 1990 IEEE International Conference on Robotics and Automation.

McGeer, T., "Powered flight, child's play, silly wheels and walking machines", 1989 IEEE International Conference on Robotics and Automation.

Pauline. P., "Construction of the Adaptive Suspension Vehicle", MIT Press, OHIO-McGraw-Hill, 1989.

Payton, D, "An Architecture for Reflexive Autonomous Vehicle Control", IEEE Computer Society Press, 1986 IEEE International Conference on Robotics and Automation, Vol3, April 7-10, 1986, P 1838-1845

Raibert. M. H., Brown. H. B., and Murthy S. S., "{3-D} balance using {2-D} algorithms?" M. Brady and R. Paul, editors, Robotics Research, the First International Symposium, pages 279--301. MIT Press, Cambridge, Ma., 1984.

Raibert M.H. and Hodgins J.K., "Biped gymnastics", The International Journal of Robotics Research, 9(2):115--132, 1990.

Raibert. M.H., "Legged Robots". Communications of the ACM, vol. 29, no. 6, pp 499-514, 1986.

Raibert M.H., "Legged robots". Communications of the ACM, 29(6):499--514, 1986.

- Raibert. M.H., "Legged Robots That Balance". MIT Press, Cambridge, Ma. 1986
- Rapp, P., et al. (1985) "Experimental Studies of Chaotic neural Behavior: Cellular Activity and Electroencephalographic Signals". Lecture Notes in Biomathematics 66:175-205.
- Rietman, Edward, "Creating Artificial Life", Windcrest, McGraw-Hill, ISBN 0-8306-4150-5, 1993.
- Rietman, Edward, "Genus Redux", Windcrest, McGraw-Hill, ISBN 0-8306-4503-9, 1994.
- Rosenblatt, J, Payton, D, "A Fine-Grained Alternative to the Subsumption Architecture for Mobile Robot Control", Proceedings of the International Joint Conference on Neural Networks, Vol. 2, June 1989, P 317-323.
- Shuji Kajita and Kazuo Tani, "Study of dynamic biped locomotion on rugged terrain - derivation and application of the linear inverted pendulum mode," 1991 IEEE International Conference on Robotics and Automation.
- Shuji Kajita, Tomio Yamaura, and Akira Kobayashi, "Dynamic walking control of a biped robot along a potential energy conserving orbit," IEEE Transactions on Robotics and Automation, 8(4):431--438, 1992.
- Skarda, C. A. and Freeman, W. J. (1987) "How Brains Make Chaos in Order to Make Sense of the World". Behavioral and Brain Sciences 10: 161-195.
- Slotine, J.E. et. al.: "Applied Nonlinear Control", Prentice-Hall, New Jersey, 1991.
- Steels, L. (1994)"The Artificial Life Roots of Artificial Intelligence"., Artificial Life Journal, Vol 1, 1. MIT Press, Cambridge.
- Tilden, M. W., "The Evolution of Functional Robo-Ecologies", ARS Electronica 93 Proceedings, LINZ Press. 1993, P195-200.
- Van Der Maas, H. L. J., Verschure, P. F. M. J. and Molenaar, P. C. M. (1990) A Note on Chaotic Behavior in Simple Neural Networks. Neural Networks 3: 119-122.
- Yamauchi, B, "JUGGLER: Real-Time Sensorimotor Control Using Independent Agents", Image Understanding and Machine Vision, Optical Society of America Technical Digest, June 1989.
- M. Vukobratovic, B. Borovac, D. Surla, and D. Stokic, "Biped Locomotion: Dynamics, Stability, Control, and Application", volume 7 of Scientific Fundamentals of Robotics, Springer-Verlag, 1990.