

BUILDING ARTIFICIAL BRAINS:

ATR, Hugo de Garis, and the Robokoneko Project

Sean McLennan

**Dr. Douglas Hofstadter
Q700: AI Hope and Hype
March 23, 1999**

1. INTRODUCTION

This paper is an overview of an enormous undertaking, both in terms of scale and implication. The “Brain Builder Group” of ATR Labs in Kyoto, Japan, headed by Hugo de Garis, is currently working on two highly related, but somewhat independent projects: the “CAM-Brain Project” and the “Robokoneko Project”. The CAM-Brain Project is primarily concerned with developing “evolvable hardware” capable of “growing” and maintaining a functioning artificial brain sufficiently powerful for real-time applications. On the other hand, the Robokoneko Project (“robokoneko” is Japanese for “robot kitten”) aims at developing a physical kitten-like robot for the artificial brain to control. Since both are being pursued in parallel by the same group, for convenience, I will henceforth refer to them as though they were a single project unless an explicit distinction is being made.

Although the project is primarily centered at ATR, much of the hardware development has been contracted out to a company named “Genobyte Inc.” (in Boulder, Colorado) under the supervision of Michael Korkin. In addition, there are around 100 other researchers worldwide involved with various tasks relating to Robokoneko and the CAM-Brain.

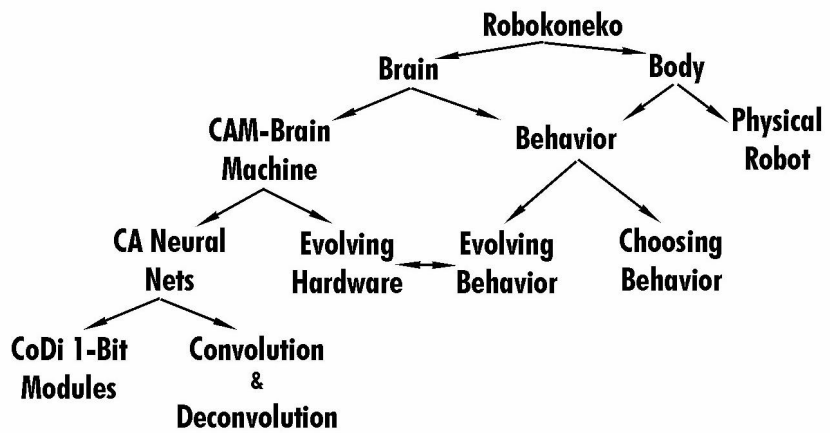
The Brain Builder Group has been underway for approximately 6 years and its originally stated goal was to have a functioning artificial brain of 1 billion artificial neurons by the year 2001. Although the artificial brain currently being developed will only have 40 million neurons, the Group is not far behind in its projected time line. They still hope that their robotic kitten will be completed by 2001.

Projects of this scope obviously address issues in a wide range of areas; section 2 provides a summary of the various aspects of the project and their challenges, successes, and current states. Section 3 discusses some of the implicit assumptions / philosophical approaches taken by the researchers and the merits and problems of the project. Finally, section 4 concludes.

2. PROJECT OVERVIEW

Perhaps it would be best to begin with a visual representation of the areas that will be discussed and their relationships as I see them.

As is clear in figure 1, the nature of this project inevitably involves an incredibly broad range of issues. In order to provide a complete picture, it is necessary to touch on all these areas; however, not every



(1). Scope of the Robokoneko Project

one will be covered in detail. Most important for this discussion are the CAM-Brain Machine, and the Evolution of Hardware and Behavior since these are the aspects most fundamental to topics in AI.

2.1. THE CAM-BRAIN MACHINE

The Brain Builder Group defines “artificial brain” as “assemblages of tens of thousands (and higher magnitudes) of evolved neural net modules” (de Garis et al. ms(a): 2). To build and house such a brain,

the Group developed the “CAM-Brain Machine” (CBM), a sophisticated piece of hardware built by Genobyte Inc. and just recently completed in March, 1999.

“CAM” stands for Cellular Automata Machine, so named because it is specifically designed to rapidly update cellular automata (CAs) which form the basic structure of the neural networks in the artificial brain (see section 2.1.1.). The original aim was to support an artificial brain of a billion neurons, one order of magnitude fewer than the human brain. In actuality, the CBM falls short of this goal, but is still capable of implementing more than 32,000 neural net modules that can each have more than 1000 neurons. Thus, a total number of approximately 40 million neurons is possible. To keep things in perspective, it is important to remember that typical neural networks as they have been implemented so far use only tens or hundreds of neurons.

The primary consideration involved in the development of the CBM has been speed. For brain-building to be plausible, tens of thousands of modules must be created within a reasonable amount of time and the completed brain must be updated fast enough to control a robot in real-time. In order to obtain the speed required, two major architectural characteristics were implemented: 1) the CoDi 1-bit neural network model (discussed in section 2.1.1.1) and 2) hardware instantiation of evolution and updating (discussed below and in section 2.1.2).

The CBM consists of 72 Xilinx XC6264 Field Programmable Gate Array (FPGA) chips. The hardware details involved are far beyond the scope of this paper (de Garis, 1997 provides an excellent overview) but basically the chips are comprised of hundreds of thousands of logic gates whose connections can be completely “rewired” by using an electromagnetic field that simultaneously reconfigures the entire chip. This reconfiguration can take place within nanoseconds, literally making it possible to have an

entirely different computer architecture with each clock tick. This is enormously beneficial for Genetic Algorithms (GAs) which are used to evolve the brain modules — the “chromosomes” of the module population to be tested, the fitness function by which they are judged, and the processes of mutation and crossover that produce more fit “off-spring” can be performed entirely *in the hardware*. Consequently, the process of evolution is performed at electronic speeds and a run of a population of 30-100 modules for 200-600 generations takes around 1 second. By comparison, MIT’s CAM-8 machine that was used before the availability of the CBM takes up to 69 minutes to perform a similar run. It has also been estimated that the CBM would be equivalent to 1,000-10,000 Pentium II 400MHz computers performing a similar task.

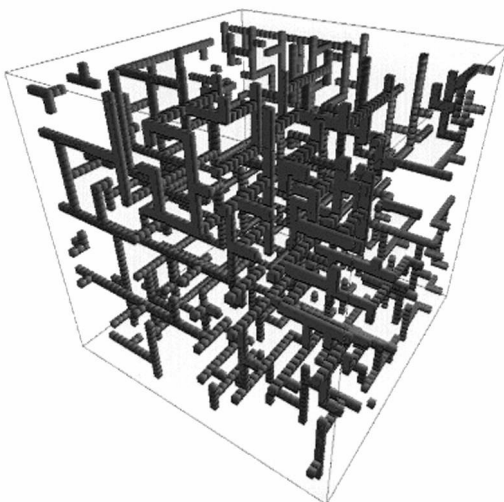
When in use, the CBM has two primary “modes”: “Evolution Mode” and “Run Mode”. As one might surmise, the process of evolution utilizes the “Evolution Mode” and the “Run Mode” is used when the brain is actually functioning. The CBM does not stand alone; it runs in conjunction with a host computer (Pentium-II 450 Mhz). The host computer communicates with the CBM via a PCI connection that provides a user interface and a few required software functions.

A summary of the technical specifications of the CBM is provided in figure 2.

Cellular Automata Update Rate (max.)	152 billion cells / sec
Cellular Automata Update Rate (min.)	114 billion cells / sec
Number of Supported Cellular Automata Cells (Max.)	453 million
Number of Supported Neurons (max., per module)	1,152
Number of Supported Neurons (max., per brain)	37,748,736
Number of Supported Neural Modules	32,768
Information Flow Rate, Neuronal Level (max.)	13.5 GB / sec
Information Flow Rate, Dendrite Level (estimated average)	40.8 GB / sec
Information Flow Rate, Intermodular Level (max.)	74 MB / sec
Number of FPGA's	72
Number of FPGA Reconfigurable Function Units	1,179,648
Phenotype/Genotype Memory	1.18 Gbytes (16MB / FPGA)
Chromosome Length	91,008 bits
Power Consumption	1 KWatt (5V, 200A)

(2). Technical Specifications of the CBM (de Garis et al. ms(a): 12)

2.1.1. CA NEURAL MODULES



As stated above, an artificial brain in the CBM can have more than 32,000 interconnected modules. They are based on a 3D cellular automata model of neural networks and are implemented within a virtual 24*24*24 cube of cells (13,824 cells in all). Figure 3 is a visual representation of a single module; the dark squares represent the neural structures, while white space are unused cells. Note that this is a virtual

(3). 3D CA Neural Module

representation — physically in the CBM, this cell space used by a single module exists within *all* 72 of the FPGA chips, each of which holds 192 cells. The chips themselves are arranged in a 6*4*3 interconnected array in the CBM. Since all the chips are used to hold the pattern of one module, it is therefore necessary

to reconfigure the entire machine for each module in order to complete an update of the entire brain. Amazingly, because of the parallel nature of the hardware (all 13,824 cells are updated simultaneously) and because the chips have a “dual-buffered” structure (allowing another module to be configured at the same time the current module is being run), the CBM is efficient enough to complete the updating of all 32,000 modules 10-20 times per second.

Of course, as in a real brain, the modules must also be interconnected and capable of passing in formation to and from each other. Each module can receive input from up to 188 other modules (or external stimulus from the robot), and send its output to up to all 32,768 modules (or external robot controls).¹ These connections are virtual and signals are stored in separate “Module Interconnection Memory” between module updates in the CBM.

2.1.1.1 CoDi-1 Bit Neural Network Model

“CoDi-1 Bit” is the name given to the CA-based neural net model that is implemented in the CBM. “CoDi” is short for “Collect and Distribute”, reflecting the nature of its neural signals, and “1 Bit” represents the signal size. Such a small signal size is uncommon for neural networks; indeed, the original design implemented an 8 bit signal. However, it became clear that to take advantage of the processing speed offered by the FPGA chips — i.e. evolution in hardware instead of software — a much simpler model was required. Thus, a 1 bit signal was dictated by hardware considerations.

There was some concern as to whether a 1 bit neural network model would be of sufficient complexity to perform adequately. To assess the viability of the model, several tests were run and modules

¹There is no explicitly stated reason for these numbers; indeed, the number of stated inputs fluctuates throughout the available documentation. 188 is the most recent value as far as I can tell.

capable of quite sophisticated functions were evolved; an XOR module, a timer module, a pattern detector, a Hubel-Wiesel Line Motion Detector, and a switchable dual function module all performed excellently (de Garis et al. ms(b)).

It is important to note that this model of neural networks does not learn. Networks that perform the desired function are developed solely through GA evolution. When a fit module is discovered, the type and orientation of each cell is stored as a 91K chromosome so that it can be recalled later when the brain is functioning in Run Mode.

Despite its signal simplicity, the model itself is quite entailed. The hand-crafted CA rules alone number well into the thousands. The reader is referred to Gers et al., 1997 for a detailed explanation of the workings of CoDi-1 Bit networks.

2.1.1.2 SPIKE INTERVAL INFORMATION CODING (SIIC)

A major issue in using a single-bit neural network model is representation. The input and output to and from the modules are binary pulse streams that de Garis and his team call “spiketrains”. At each clock tick, either a 1 — a “spike” — or a 0 is produced (or input) but what those spikes actually *mean* in terms of controlling a robot or recognizing motion etc. is not a simple matter.

There are several possible methods of interpreting spiketrains. A typical frequency based approach (i.e. parsing the signals into windows of “n” clock cycles) would cause an unacceptable reduction in processing speed. An ideal representation would interpret an integer value at every cycle. “Spike Interval Information Coding” (SIIC) is such a representation. The details of this encoding/decoding method are extremely entailed and worthy of study in their own right, but essentially, the method uses a static filter (also

evolved using GAs) to convert spiketrains to and from analog signals. A complete account can be found in Korkin et al, 1998.

2.1.2. EVOLVING HARDWARE

There are two “phases” in the CoDi 1-Bit model: the “Growth Phase” (used to grow a module) and the “Neural Signaling Phase” (used to propagate signals through an existing module). The process of evolution in the CBM makes use of both.

Each cell can be blank or one of three types: neuron, axon, or dendrite. Since the network is implemented in 3D space (24*24*24 cells), each cell has 6 possible neighbors; neuron and dendrite cells each have five inputs and one output, axon cells are the opposite and have one input and five outputs. At the start of the Growth Phase, the cell space is randomly seeded with neurons. The neurons send “growth signals” to their neighbors which in turn propagate other growth signals according to hand-coded CA rules. Synapses are created when the growing trails intersect each other.


After the module has been grown, it enters the Neural Signaling Phase and its output is evaluated on how well it performs the targeted task. Fitness evaluation is also completely instantiated in a hardware unit which consists of the input spiketrains, target spiketrains, and the evaluator (all of which are humanly designed and hand-coded). At the end of the signaling phase, the module’s fitness is instantly provided to the Genetic Algorithm Unit.

The selection of the best modules is performed in software on the host computer, but the generation of offspring (by crossover and mutation) is performed in hardware in the CBM to take advantage of

electronic speeds. Typically, a suitable module will be evolved from a population of 30-100 modules in 200-600 generations in about one second.

2.2. ROBOKONEKO

The primary interest of the Brain Building Group is to create brains — not robots. However, they recognized that their brain would need to *control something* and they settled on a robotic kitten. The Group quite freely admits that the motivation of choosing this particular form was to draw media attention. If successful, the creation of a cute, frisky robotic pet will be vivid proof of the efficacy of the CAM-Brain Project and its capabilities will be apparent to scholars and lay-persons alike. The Robokoneko model is depicted in figure 4.



(4). Robokoneko

2.2.1. BODY

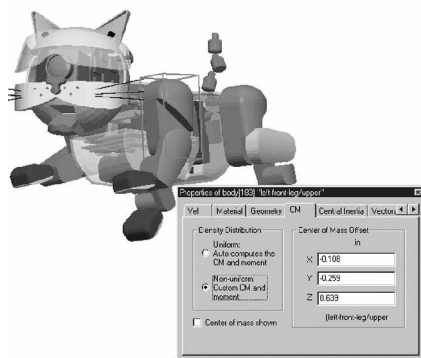
At present, Robokoneko does not exist in a physical form; building the actual robot will depend on the outcome of simulations that are currently being run. The Group does not want to incur the cost of a real world robot (estimated to be in the tens of thousands of dollars) until certain obstacles have been

overcome. Namely, the viability of evolving motion control is a major issue (discussed below in section 2.2.2.1). However, the body has been designed and when the Group is ready, the robot's construction will be contracted to Genobyte Inc., the company headed by Michael Korkin that built the CBM.


The projected weight of the kitten is about 3 kilos and its length, about 25 cms. Obviously, Robokoneko is not “brains-on-board” since its brain is contained in the CBM — it will be controlled via radio. The torso is comprised of two parts joined with 2 degrees of freedom (DoF). The back legs have 1 DoF at the ankle and knee, and 2 DoF at the hip. The front legs have 1 DoF at the knee and 2 at the hip. All 4 feet are spring loaded. It has also 3 DoF at the neck, 2 at the tail, and 1 at the mouth. All told, there will be more than 20 motors in Robokoneko — one for each DoF at each position. Robokoneko's senses include a single CCD camera eye, two microphone ears, a gyroscope, and various touch sensors. Additionally, a battery pack, radio receiver and antenna, sound chip, and other control and sensory pre-processing hardware (image digitization for example) will be on board. It is expected that Robokoneko will be primarily made out of light-weight plastics and will be able to function for about 20-30 minutes on battery power. All of the necessary components to build the robot exist on the market today.

Although the real robot has not yet been created, Robokoneko has been simulated with a sophisticated software package called “Working Model 3D” (WM3D) produced commercially by Knowledge Revolution Inc. WM3D allows the user to not only design machines, but simulate materials, forces, and working conditions. Some screen captures of the Robokoneko simulation in WM3D are presented in figures 5 to 7.

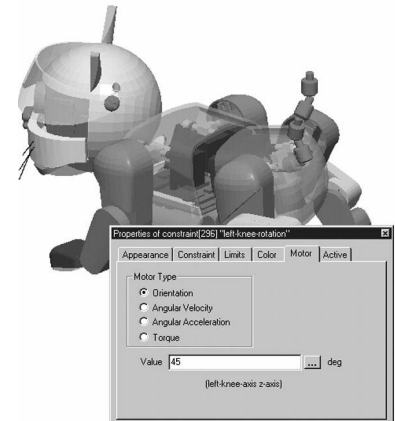
(5). Robokoneko - Materials



(6). Robokoneko - Interior



(7). Robokoneko - Motors




2.2.2. BEHAVIOR

Robokoneko will not begin with a full 32,000 module brain — a much smaller set will be attempted first. The Brain Building Group seems to be approaching the robot's control as 1 behavior = 1 neural module, although the definition of "behavior" has not been explicitly outlined. Some of the "behaviors" they intend to evolve include: walk straight, turn left, turn right, run straight, sit, jump, pounce, play with a ball, strike out with a paw, turn head to follow an object, meow (with various emotions), and show a short attention span (Agah et al, ms). Clearly, here, "behavior" is being treated as a broad concept without discrete units; presumably "play with a ball" would entail "strike out with a paw", both of which are quite dramatically different from "show a short attention span".

However, despite this "fuzziness of definition", modules evolved so far seem to be successful in producing complicated sequences of motor control (each of which might be considered several "behaviors" in their own right) in real-time simulations. It seems that each module is complex enough to support a great deal of behavioral information.

2.2.2.1 EVOLVING BEHAVIOR



(8). Block Model of Robokoneko

The evolution of the modules used for the robot's motion pose a unique problem that requires a violation of the philosophy of the project: i.e. they can not rely on hardware instantiated evolution to develop real world behavioral control. The problem is speed and accuracy; if motion control was evolved mechanically, i.e. each module was tested using a real robot, the performance

of each module evaluated in the GA would have to be judged by a human observing the robot. Considering that the evolution of a fit module requires evaluations of approximately 60,000 other modules, it becomes clear that the time needed to evolve an entire brain would render the project unfeasible. In addition, casual human observation would not be able to distinguish the average 1% increase in fitness typically observed between generations.

Thus, an interface between WM3D and the CBM was developed (a software package called "RobotStim" by Genobyte Inc.) so that modules could be evolved using the simulated Robokoneko, with the justification that the results of the simulated evolution could be used as a starting point to further fine-tune the behavior of the real robot. Unfortunately, because WM3D runs on a conventional PC, it too does not run fast enough for evolution using the full Robokoneko model. Consequently, the group felt obliged to use a simpler block model (shown in figure 8.), and to hand-encode the starting population's chromosomes with approximate starting values in order to speed up evolution. The justification is the same: the fittest modules from the simulation of the simple model can be used as a starting point for the more complex model. Therefore the process of developing motion control modules will go through a total of four

steps: human approximation of starting values, simulation using the block-Robokoneko model, simulation using the full-Robokoneko model, and finally, real world fine-tuning.

As of this writing (March, 1999), motion modules have only been evolved at the block-model level so it remains to be seen if this approach will be feasible. Behaviors that have been developed so far include “walk”, “veer left”, “bow”, “arch back”, “breast stroke”, and “stretch”; the success of these modules demonstrate the validity of at least the initial stages of the process.

3. DISCUSSION

Discussion of the Brain Building Project will be divided in to 3 sections: 1) the researcher’s philosophical approaches, 2) the project’s merits, and 3) the project’s drawbacks.

3.1. PHILOSOPHICAL APPROACHES OF THE BRAIN BUILDING GROUP

Embodiment

There has been a strong movement in the field of Cognitive Science and Robotics towards what has been dubbed “embodiment”; i.e. control systems that function in the real world. (The work of Rodney Brooks (Brooks, 1997), one of the principle’s firmest proponents, is a good example.) Clearly this is the approach that has been taken at ATR from the beginning; they have always intended for their artificial brain to be used for real-world, real-time applications.

Emergence

“Emergence” is a blanket term covering a wide variety of phenomena from the construction of insect nests to human cognition. In spirit, it means that simple rules and patterns on one level, when allowed to interact in great numbers, will result in unpredictable and complex behaviors and patterns on a higher level.

The Robokoneko project relies on emergence in two areas — the use of cellular automata neural nets and the use of genetic algorithms to evolve them. In most cases, the brain modules evolved in the CBM begin with a random state and are “grown” via a set of rules that look at a cell and its immediate, local environment to determine what changes will occur in the next time step. The result of these local interactions is a structure of neuronal trees of such complexity that it is difficult for humans to envision, let alone design.

Of course, randomly generated CA networks are not likely to perform a specified task at all, let alone adequately, hence the use of GAs. Similarly, simple rules (a user defined fitness function, random processes of mating and mutation, and selection of the fittest) applied repeatedly over a substantial population for several generations result in an extremely effective search of “solution space” producing a good, if not perfect, network structure. Again, a structure that would be nearly impossible for a human to design.

3.2. MERITS

By far the strongest evidence of the merit of the CAM-Brain Project is what they have thus far produced. The CBM itself — a computer capable of evolving its own architecture — is a monumental achievement whose applications are only limited by our inability to conceive of such a radically different form of computation. Even if Robokoneko proves not to be a viable project, the CBM and similarly designed

computers would have value independent of their use in building artificial brains. As with all computer technology today, the CBM is already out of date despite the fact it's not yet a month old; Xilinx has already unveiled their latest FPGA chip which contains a million logic gates — more than 10 times the number of gates on the chips in the CBM.

In addition to the architecture of the CBM, the modules that have so far been produced by the CBM have proven to be successful in accomplishing sophisticated functions. The CBM could be used to efficiently design more conventional neural networks for applications to be implemented in more conventional devices. Even if designing the architecture of 32,000 interconnected modules proves to be impossible, the CBM has demonstrated a marked improvement in the development of smaller-scale networks. SIIC encoding, too, provides a significant tool whose possibilities have not yet even been considered.

This project has successfully integrated a wide variety of related areas of study demonstrating not only the value of an interdisciplinary approach in brain building, but also serving as an example of how other problems might be tackled. Cellular automata, for example, are surely fascinating but they have not been taken very seriously because their practical applications are not immediately clear. Similarly, the applications of neural networks have been limited by scale, and the applications of genetic algorithms by speed. The CAM-Brain Project has finally brought together these three somewhat related fields and shown how they can complement each other. Pedagogically, the project has a lot to offer the cognitive science community.

Finally, the Brain Builder Group shows merit in that they are currently not far off their projected time line. If they continue to progress at the current rate, it is conceivable they will reach their stated goal of producing Robokoneko by 2001. Certainly, the amount of work yet to be done is intimidating, but the Group's past success allows one to be optimistic.

3.3. DRAWBACKS

The approach of "evolutionary programming" or "evolutionary engineering" such as implemented by the researchers at ATR, has a rather large drawback — one that has been brought up many times and is acknowledged by the Brain Building Group themselves. Although the final product performs as desired, even its creators have no idea *how* it does what it does. The structure and dynamics of the system remain a largely unanalyzable "black box". Speaking of the successful evolution of a Hubel-Wiesel Line Motion Detector, de Garis et al. (ms(a): 17) state:

"Of course, we have no idea how the circuit does what it does. This is the great strength of 'evolutionary engineering'. Evolved circuits can achieve performance levels beyond what human engineers can achieve with traditional top-down design techniques, i.e. attain superior engineering performance levels, but the price is that one loses scientific understanding, due to the overwhelming structural and dynamical complexity of these CoDi circuits. Thus 'evolutionary engineering' can provide superior engineering, but inferior science. It is a trade-off. In practice, once EEs can generate tens of thousands, even millions of modules, only a few die-hard analysts will want to know how an individual module functions. For the most part, noone [sic] will care how a particular module amongst millions actually does its thing."

Indeed, this statement harkens back to one of the Brain Builder Group's stated fundamental assumptions (de Garis. 1994: 2):

"...hyper complex systems (such as biological brains or embryos) will probably have to be built using an evolutionary approach rather than using human design."

Thus, it seems (at least according to Hugo de Garis) Cognitive Science is doomed to have little understanding of the fundamental processes that underlie cognition despite being able to produce it. I feel this claim is too strong.

It may be true that the average person will not care how artificial brains do what they do (how many people truly care how their TV works?) but de Garis overlooks an important fact. If brain building proves successful, it will be possible to use the CBM as a tool for studying precisely the phenomena it exhibits. A great deal of science begins “in the middle” with the ability to reproduce a phenomenon without being able to completely understand it. Case in point: the development of GAs has (in retrospect) provided some insights into real-world evolution (Levy, 1992: 215-230).

The advantage is that with the CBM, we know the entire structure of each module down to minute detail and it would be theoretically possible to record every activation and every signal produced as the brain functions. Although eyeballing that data would hardly be plausible, it is not difficult to imagine that “a die-hard analyst” could come up with a ingenious method of analyzing it computationally.

Another possible drawback to the Group’s approach regards behavior. It seems that there will be a lot of functional overlap between the modules — “playing with a ball” involves “walk forward” which in turn involves “lift left foreleg” which in turn entails “bend knee”. Intuitively, this seems like an inefficient design; our immediate instinct is to limit redundancy and to reduce complex behavior into sub-tasks that are then interrelated and ordered.

To now (perhaps perversely) come to the Group’s defense, this approach may not be as invalid as it seems. Recent developmental research (Thelan, 1995) suggests that the learning of motor control is highly task-dependent and what is learned in one task does not immediately transfer to another. Theoretical

issues aside, what is evident is that Robokoneko can perform well, even if only in a simplified simulation; it may be that the final robot will behave as desired despite an aesthetically displeasing design.

Clearly, there is still an immense amount of work yet to be done before Robokoneko is fully realized. One major issue that has not been adequately addressed in the work presented so far is the interconnectivity of the brain modules themselves. Hugo de Garis and the group feel that the development of a single module of 13,824 cells and 1000 neurons is far beyond human capabilities, yet continue to maintain that the modules themselves — up to 32,000 of them — will be put together according to a human-designed architecture. Not only is this a task of several magnitudes greater than designing a single module, it also seems to contradict their basic premise of relying on evolution. To be fair, until now the Group's primary concern has been simply producing single modules and now with that task accomplished, they are only now turning to assembly issues. It is, however, odd that no one seems to have considered how it might be possible to apply the same evolutionary principles to the next level of the project.

It is true that the group will not be attempting to assemble a full 32,000 module brain immediately; they will begin with much smaller structures in the tens and hundreds. Perhaps what they learn in those attempts will lead to insights that will facilitate the herculean task of designing a complete brain architecture.

One possibility might also eliminate another drawback to the Brain Builders' approach; implementing learning on some level. de Garis et al (ms(a): 18) have already speculated on the possibility of learning in the CoDi 1-Bit model via reverberating internal signals. In addition, there seems to be no immediately apparent reason why modules could not be configured to interact directly with the CMB to direct the further evolution / re-evolution of modules in real-time based on interactions between the robot

and its environment. The dynamic nature of the CBM's FPGA chips leave open realms of possibilities not even imagined, let alone explored.

Although these issues are serious ones, the inventiveness of the Brain Builder Group has already proven more than adequate against other just as serious obstacles in the past. Considering what has already been accomplished, one is inclined to allow them the benefit of the doubt at least as long as they seem to be advancing towards their goal.

4. CONCLUSION

This paper presented an overview of a monumental undertaking — the building / growing / evolving of a fully functioning artificial brain intended to control a real-world robot in real-time. Although the project is not yet complete, the team of researchers in ATR's Brain Building Group have achieved several significant steps towards their goal of producing a playful, robotic kitten by 2001. Those successes include: creating the CAM-Brain Machine, a computer capable of evolving CA Neural Network modules in about a second; developing a unary encoding method that optimizes processing speed; evolving neural network modules capable of performing sophisticated functions; and evolving neural network modules for real-time motion control of a simulated robot.

What remains to be accomplished is to complete evolving the required tens and hundreds more behavioral, sensory, and decision-making modules required, to design the module-level brain architecture that will connect them, and to build the physical robot. This may prove to be the most difficult task, but even if brain building using this approach is not successful, the project has paved the way for a great deal more exploration in any number of possible directions.

REFERENCES

- Agah, Arvin, Hugo de Garis, Michael Korkin, Katsunori Shimohara, and Sung-Bae Cho. ms. "Architectural and Functional Specifications for a Robot Kitten 'Robokoneko' to be Controlled by a 10,000 Evolved Neural Net Module Artificial Brain."
- Brooks, R. A. 1997. Intelligence without Representation. In J. Haugeland (Ed.) *Mind Design II*. pp. 395-420. MIT Press: Cambridge, MA.
- de Garis, Hugo, Michael Korkin, Katsunori Shimohara. ms(a). "ATR's 'CAM-Brain Machine' (CBM) and Artificial Brains: An FPGA Based Hardware Tool which Evolves a Neural Ner Circuit Module in a Second and Updates a 40 Million Neuron Artificial Brain in Real Time." submitted to: *Journal of VLSI Signal Processing Systems (JVSPS), Special Issue on Custom Computing Technology*.
- de Garis, Hugo, Andrzej Buller, Michael Korkin, Felix Gers, Eiji Nawa, Michael Hough. ms(b). "ATR's Artificial Brain ('CAM-Brain') Project: A Sample of What Individual 'CoDi-1Bit' Model Evolved Neural Net Modules Can Do with Digital and Analog I/O."
- de Garis, Hugo, Nikolai Petroff, Michael Korkin, Gary Fehr, and Eiji Nawa. ms(c). "A 40 Million Neuron Artificial Brain for an Adaptive Robot Kitten 'Robokoneko'". submitted to: *Connection Science Journal, Special Issue on Adaptive Robots*.
- de Garis, Hugo. 1993. "Evolvable Hardware: The Genetic Programming of Darwin Machines." International Conference on Artificial Neural Networks and Genetic Algorithms. Innsbruck, Austria.
- de Garis, Hugo. 1994. "CAM-Brain: Growing an Artificial Brain with a Million Neural Net Modules Inside a Trillion Cell Cellular Automata Machine." *Journal of the Society of Instrument and Control Engineers of Japan (SICE)*. Vol. 33, No. 2. Society of Instrument and Control Engineers of Japan.
- de Garis, Hugo. 1997. "Evolvable Hardware: Principles and Practice." *Communications of the Association for Computer Machinery (CACM Journal)*.
- Korkin, Michael, Norberto Eiji Nawa, and Hugo de Garis. 1998. "A 'Spike Interval Information Coding' Representation for ATR's CAM-Brain Machine (CBM)." Int. Conf. on Evolvable Systems. Lausanne, Switzerland.
- Gers, Felix, Hugo de Garis, and Michael Korkin. 1997. "CoDi-1Bit: A Simplified Cellular Automata Based Neuron Model." AE97 Artificial Evolution Conference. Nimes, France.
- Levy, Steven. 1992. *Artificial Life*. Vintage Books: New York.

Thelen, Esther. 1995. "Time-scale dynamics and the development of an embodied cognition." In Port and van Gelder (eds). *Mind as Motion: Explorations in the Dynamics of Cognition*. Bradford Books, MIT Press: Cambridge, MA. pp. 69-100.

WEB REFERENCES

Hugo de Garis and ATR (including all CAM-Brain and Robokoneko project papers listed above):
<http://www.hip.atr.co.jp/~degaris/>

Genobyte Inc.: <http://www.genobyte.com/>

Xilinx Inc.: <http://www.xilinx.com/>

Working Model 3d: <http://www.workingmodel.com/>

Sean McLennan: <http://www.norwestcontols.com/Sean/>