The China-Brain Project

Report on the First Six Months

Prof. Dr. Hugo de GARIS, Dean Prof. Dr. ZHOU Changle, Prof. Dr. SHI Xiaodong, Dr. Ben GOERTZEL, Prof. PAN Wei, Prof. MIAO Kehua, Prof. Dr. ZHOU Jianyang, Dr. JIANG Min, Prof. ZHEN Lingxiang, Prof. Dr. WU Qinfang, Prof. Dr. SHI Minghui, LIAN Ruiting, CHEN Ying

Artificial Brain Lab
Cognitive Science Department
School of Information Science and Technology
Xiamen University, Xiamen, China

profhugodegaris@yahoo.com

Abstract

The "China Brain Project" is a 4 year (2008-2011), 10.5 million RMB research project to build China's first artificial brain, which will consist of 10,000-50,000 neural net modules which are evolved rapidly in special FPGA hardware, downloaded one by one into a PC or supercomputer, and then connected according to the designs of human "BAs" (brain architects) to build an artificial brain with thousands of pattern detectors to control the hundreds of behaviors of a two legged robot.

1. Introduction

The "China Brain Project", based at Xiamen University, is a 4 year (2008-2011), 10.5 million RMB, 20 person, research project to design and build China's first artificial brain (AB). An artificial brain is defined here to be a "network of (evolved neural) networks", where each neural net(work) module performs some simple task (e.g. recognizes someone's face, lifts an arm of a robot, etc), somewhat similar to Minsky's idea of a "society of mind" [1], i.e. where large numbers of unintelligent "agents" link up to create an intelligent "society of agents". 10,000s of these neural net modules are evolved rapidly, one at a time, in special (FPGA based) hardware and then downloaded into a PC (or more probably, a supercomputer PC cluster). Human "BAs" (brain architects) then connect these evolved modules according to their human designs to architect artificial brains. Special software, called IMSI (see section 5) is used to specify these connections, and to perform the neural signaling of the whole brain (in real time). The main aim of this research project is to show that using this (evolutionary engineering) approach to brain building is realistic, by simply building one and show that it can have thousands of pattern recognizers, and hundreds of motions that are switched between, depending on external and internal stimuli. This project already has (Fujian) province "key lab" financial support. It is hoped, in three years, that it will win "key *state* (i.e. federal) lab" status. In 5-10 years, it is hoped that China will establish a "CABA" (Chinese Artificial Brain Administration), consisting of thousands of scientists and engineers, to build national brains to fuel the home robot industry (which may become the worlds largest) (See section 11.)

There are about 20 people (8 professors) involved in this project, divided into specialist teams, i.e.

a) "Vision team" (who evolve pattern recognition modules and create vision system architectures).

b) "Robot team" (who program the NAO robot [2] to perform the many (hundreds of) behaviors that the robot is to perform).

c) "Hardware Acceleration team" (who program the FPGA electronic boards we use to evolve neural net modules as quickly as possible).

d) “Supercomputer team” (who port the Parcone and IMSI code to a supercomputer, to accelerate the neural evolution and signaling of the artificial brain).

e) “Language team” (who give the robot language capabilities, i.e. speech, listening, understanding)

f) “Consciousness team” (who aim to give the NAO robot some degree of self awareness).

By the end of this 4 year project, we hope to be able to show the world an artificial brain, consisting of 10,000s of evolved neural net modules that control the 100s of behaviors of a NAO robot (and probably a more sophisticated robot) that makes a casual observer feel that the robot “has a brain behind it”. It also hoped that this artificial brain project will encourage other research teams to work in this new area, as well as help establish an artificial brain industry that will stimulate the growth of the home robot industry, probably the world’s largest by 2030.

2. The “Parcone” (Partially Connected Neural Evolutionary) Neural Net Model

If one chooses to build an artificial brain based on the “evolutionary engineering” of neural net modules, then the choice of the neural net model that one uses to evolve the modules is critical, since everything else follows from it. Hence quite some thought went into its choice. We eventually decided upon a partially connected model (that we called the “Parcone” (i.e. partially connected neural evolutionary) model, since we wanted to be able to input images from digital cameras that started off as mega-pixel images, which were then compressed to 1000s to 10,000s of pixels. In earlier work, the first author [3], had always used fully connected neural networks for his neural net evolution work, but with 10,000s of pixel inputs (with one pixel per input neuron) a fully connected neuron would have an unreasonably large number of connections (i.e. hundreds of millions).

The moment one chooses a partially connected neural net model, one must then keep a list of all the neurons that each neuron connects to. This we did in the form of a hash table. See the data structures of Fig. 2. Each neuron that is connected to from a given neuron has a unique integer ID that is hashed to find its index in the hash table of the given neuron. This hash table slot contains a pointer to a struct that contains the integer ID of the neuron connected to, the weight bits of the connection, and the decimal weight value.

These connections and weight values are used to calculate the neural output signal of the given neuron. The weight bits are mutated during the evolution of the neural net, as well as the connections, by adding and cutting them randomly. The model contained many parameters that are chosen by the user, e.g. the number of input, middle, and output neurons, the number of weight-sign bits, the number of hash table slots, the population size, mutation rates, etc. These parameters were “tuned” empirically for maximum evolution speed.

3. Pattern Detection Results

Once the Parcone code was written and debugged, the first pattern recognition task we undertook was to see how well it could recognize faces. Fig. 3 shows an example of the face inputs we used. We took photos of 6 people, with 5 images of each person at different angles, as Fig. 3 shows. 3 of these images of a given person were used as the positive cases in the training set, plus 3 each of two other people, as the negative cases in the training set. The Parcone neural net was evolved to output a strong positive neural signal for the positive cases, and a strong negative signal for the negative cases. When the other positive images not seen by the evolved module were input, the outputs were strong positive, so the module generalized well.

We then automated the pattern detection so that P positive images could be used in the evolution, and N negative images. Users could then select from a large menu of images the P positive and N negative images they wanted in the evolution. We evolved shoe detectors, mobile phone detectors etc. in a similar way. When a shoe detector was presented with a face, it rejected it (i.e. it output a negative neural signal) in about 95% of cases, and vice versa. Thus we found that when a detector for object “X” was evolved, it rejected objects of type “not X”. This ability of the Parcone model will hopefully prove to be very useful for constructing the 1000s of pattern detectors for the artificial brain.

At the time of writing (Oct 2008) tests are currently underway by the vision group to evolve *motion* detectors. A moving image is input as a set of "movie frames". Once we have the Parcone model implemented in FPGA based electronics (see section 6) we hope to be able to evolve pattern detectors (whether stationary or moving) in real time (i.e. in about 1 second).

4. The NAO (Robocup Robot Standard) Robot

Fig. 1 shows the NAO robot, a product of the French company Aldebaran [2], in Paris. It is of waist height, costs about $20,000, can walk on its two legs, talk, listen, grip with a thumb and two fingers, and has one eye. It is now the Robocup robot standard, after Sony stopped supporting their Aibo robo-pet dog, which was the previous Robocup robot standard. Interestingly for our project, the NAO robot (which means "brain" in Chinese by the way (coincidence? marketing?)) comes with accompanying motion control software, called "Choregraphe" which we have chosen to use, rather than try to evolve motion control for all the 25 motors that the NAO robot possesses. We expect to have hundreds of motions for the NAO robot so that it can accomplish many tasks that the artificial brain initiates.

We are acutely conscious that no one will actually "see" the artificial brain, since it will consist of 10,000s of neural net modules hidden inside the PC or supercomputer that performs the neural signaling of the artificial brain. All that human observers will see will be the robot, that the artificial brain controls, so we are hoping that when observers watch the hundreds of behaviors of the robot, with its 1000s of pattern detectors, they will have the impression that the NAO robot "has a brain behind it" and be suitably impressed (at least enough for the funding of the research project to be continued). Later in the project, we intend building a more sophisticated robot with two eyes, and hands with better fingers, capable of real grasping, with touch sensors, etc. The project has a dedicated "robot group" who work on generating its motions, and control.

5. IMSI (Inter Module Signaling Interface)

IMSI stands for "inter module signaling interface", i.e. the software "operating system" that is used for several purposes, namely :-

a) It allows the "BAs" (brain architects, i.e. the people who decide which neural net modules to evolve (i.e. their fitness definitions, etc) and the architectures of artificial brains) to specify the connections between the modules (e.g. the output of module M2849 (which performs task "X") connects to the 2nd input of module M9361 (which performs task "Y")). Such information is stored in special look up tables (LUTs).

b) These LUTs are then used to allow the IMSI to perform the neural signaling of the whole artificial brain. When the output signal is being calculated for a particular module, it needs to know the values of the neural signals it is getting from other modules, and to which modules to send its output signal.

The IMSI calculates the output neural signal values of each module sequentially, for all modules in the artificial brain. Placing dummy weight values for about 1000 connections per module, allowed us to use a PC to determine how many such "dummy" modules could have their neural output signals calculated sequentially in "real time" (defined to be 25 output signals for every *neuron* in the artificial brain). The answer was 10,000s depending on the speed of the PC. Since we have a 10-20 PC node supercomputer cluster at our disposal, we can realistically envision building an artificial brain with several 10,000s of neural net modules.

At first, we envisioned that the artificial brain would consist solely of evolved neural net modules, interconnected appropriately to generate the functionality we desired. However the decision mentioned in section 4 (on the NAO robot) that we would use Aldebaran's "Choregraphe" software to control the motions of the robot's many behaviors, implies that the IMSI will be a hybrid of neural net modules and motion control routines written with Choregraphe.

The IMSI software will be designed and coded in November of 2008, allowing the first "microbrain", consisting of some dozen or so modules, to be designed and tested.

6. FPGA Based Parcone Module Evolution

The accelerator group is using 3 FPGA electronic boards to evolve neural net modules (based on the Parcone model) as quickly as possible, hopefully in real time (i.e. in about a second). Real time evolution will allow continuous learning of the artificial brain, so evolution speed has always been a dominant factor in our research. If these FPGA boards prove to be too slow, we may try a hybrid analog-digital approach, where the neural signaling is done using analog neurons based on Prof. Chua's "CNN" (cellular neural networks) [4], and the evolution is controlled digitally. This latter approach will demand a much higher learning curve, so will not be undertaken if the FPGA board approach proves to be sufficient.

7. The Language Component

The NAO robot is to be given language capabilities. Prof. SHI Xiaodong and Dr. Ben GOERTZEL are in charge of the "Language team". The NAO robot (actually, the artificial brain) is to be made capable of speaking, listening, and understanding spoken commands and answering spoken questions. This will involve speech to text and text to speech conversion, which will probably use standard "off the shelf" products. The research effort will be based more on language understanding, parsing, etc. The aim is to be able to give the robot spoken commands, e.g. "Go to the door". "Pick up the pencil on the table", etc. The robot should also be capable of giving verbal replies to simple questions, e.g. the question "Where is Professor X" might get an answer "Near the window".

8. The Consciousness (Self Awareness) Component

Dean Zhou Changle, (dean of the School of Information Science and Technology) at Xiamen University, is responsible for the consciousness (self awareness) component of the project. At the time of writing (Oct 2008), this component is still under consideration. The dean is keen on this component, even referring to this project as the "Conscious Robot Project".

9. Near Future Work

At the time of writing (Oct 2008), the China Brain Project is only about 6 months old, so there is not a lot to report on in terms of completed work. Now that the evolvable neural net model (Parcone) is complete and tested, the most pressing task is to put a (compressed) version of it into the FPGA boards and (hopefully) speed up the evolution of a neural net (Parcone) module so that it takes less than a second. This "real time" evolution then opens up an exciting prospect. It would allow "real time" continuous learning. For example – imagine the robot sees an object it has never seen before. All the pattern recognition circuits it has already stored in its artificial brain give weak, i.e. negative output signals. Hence the robot brain can detect that the object is not recognized, hence a new pattern detector circuit can then be learned in real time and stored.

The robot group will use the Choregraphe software to generate hundreds of different behaviors of the NAO robot.

The vision group will continue testing the Parcone model for evolving pattern detectors, e.g. detecting motion (e.g. distinguishing objects moving left from those moving right, between those moving towards the eye of the robot quickly, from those that are moving slowly, etc). There may be thousands of pattern detectors in the artificial brain by the time the project contract finishes at the end of 2011.

We hope to have integrated a language component by March of 2009 (before the AGI-09 conference), so that we can have the robot obey elementary spoken commands, e.g. "move to the door", " point to the window", "what is my name?" etc.

Also by then, the aims of the "consciousness (self awareness) component" should be better clarified and whose implementation should have begun.

10. Goals for the Next Few Years

The first year is devoted largely to *tool building* (e.g. choosing and coding the evolvable neural net model, testing its evolvability as a pattern detector, implementing the Parcone model in the FPGA boards, programming the motions of the NAO robot with the Choregraphe software, writing the IMSI code, etc). In the second year, the first artificial brain *architectures* will be created and implemented, with 10, 20, 50, 100, etc modules. The language component will be added to the brain. At regular intervals, demos will be built, to show off progress. We expect that each scaling up of the size of the artificial brain (i.e. each substantial increase in the number of modules in the brain) will raise new challenges that will have to be overcome by the creativity of the team's BAs (brain architects). At the end of 4 years, it is hoped that the artificial brain will have 10,000s of modules. Observing a robot controlled by such an artificial brain should make a casual observer feel that the robot "has a brain behind it".

11. Artificial Brain Research Policy in China

The first author thinks that the artificial brain industry will be the world's biggest by about 2030, because artificial brains will be needed to control the home robots that everyone will be prepared to spend big money on, if they become genuinely intelligent and hence useful (e.g. baby sitting the kids, taking the dog for a walk, cleaning the house, washing the dishes, reading stories, educating its owners etc). China has been catching up fast with the western countries for decades. The first author thinks that China should now aim to start leading the world (given its huge population, and its 3 times greater average economic growth rate compared to the US) by aiming to dominate the artificial brain industry. At the time of writing (Oct 2008), plans are afoot (with the support of the most powerful people in China in artificial intelligence) to attempt to persuade the Chinese Ministers of Education and of Science and Technology to undertake a long term strategy (over a 20 year time frame) to dominate the global artificial brain industry, by initially stimulating the establishment of artificial brain (and intelligence science) labs in universities across China (somewhat similar to our Artificial Brain Lab at Xiamen), then awarding "key lab" status at both province and national levels to some of these labs, the creation of a Chinese National Artificial Brain Association, and especially, within a 5 to 10 year time frame, the establishment of a "CABA" (Chinese Artificial Brain Administration), which would be a government administration (similar to America's NASA) that would employ thousands of engineers and scientists to build artificial brains for the Chinese artificial brain industry. Copying the human brain with its 100 billion neurons and quadrillion synapses will be an immense task requiring large numbers of Chinese "brain workers". It is expected that other countries will quickly copy China's lead, so that one will soon see national brain building projects in most of the leading high tech nations.

References

[1] Marvin MINSKY, "Society of Mind", Simon & Schuster, 1988.

[2] http://www.aldebaran-robotics.com/eng/index.php

[3] Hugo de Garis, in book, Artificial General Intelligence 2008 : Proceedings of the First AGI Conference (Frontiers in Artificial Intelligence and Applications) (Frontiers in Artificial Intelligence and Applications)

[4] Leon CHUA and Tamas ROSKA, "Cellular Neural Networks and Visual Computing, Foundations and Applications", Cambridge University Press, 2002.

Fig. 1 The NAO Robot Controlled by Our Artificial Brain

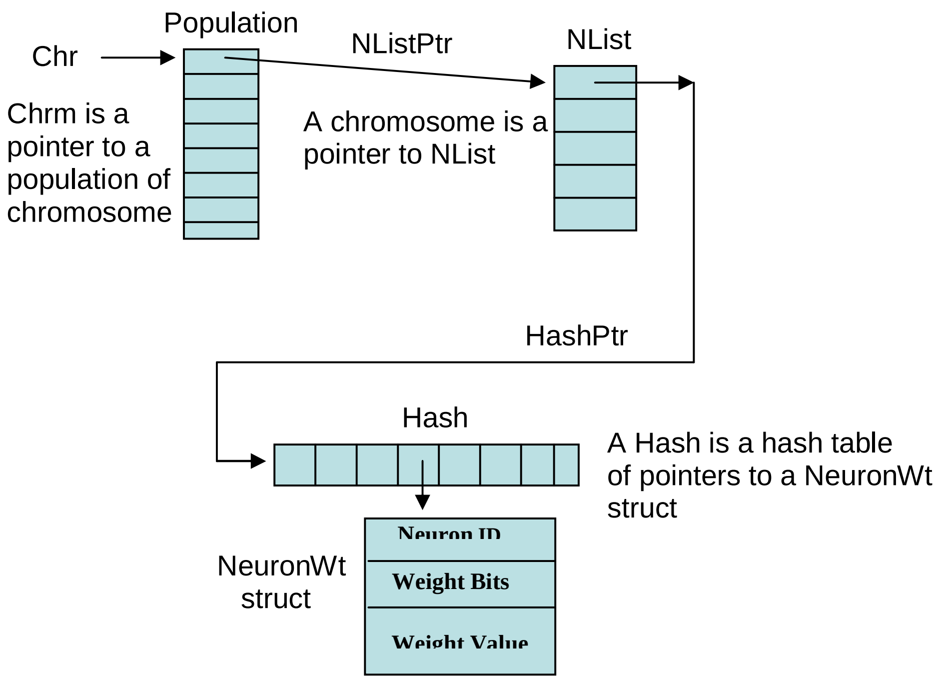

Fig. 2 Data Structures of the "Parcone" Neural Net Model

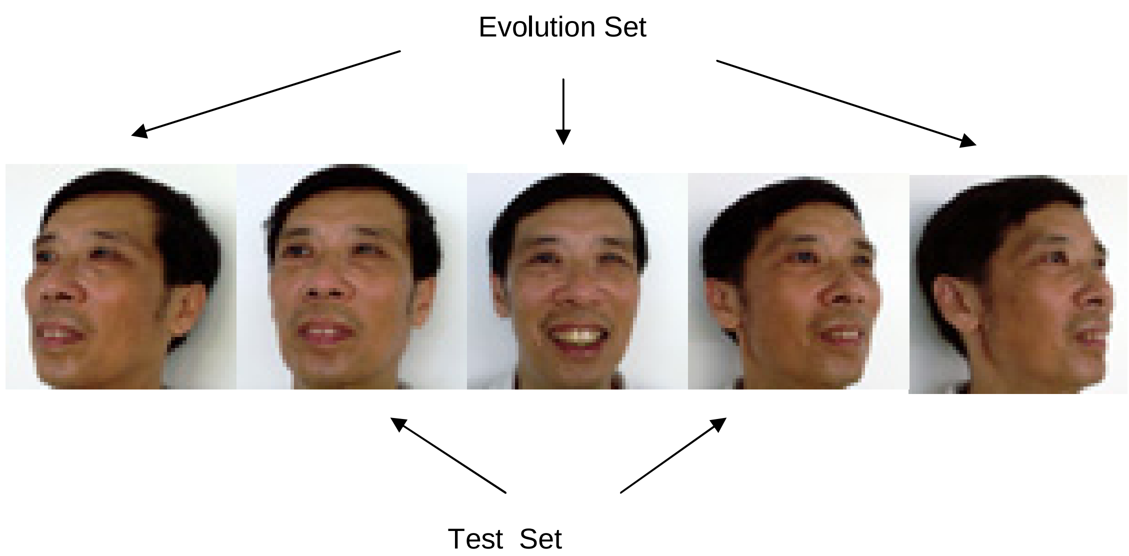

60*60 (pixel)*3 (RGB) = 10800 pixel values = 10800 input neurons to Parcone

Fig. 3 Face Images for Face Detection Evolution of Parcone