BrainOS – A theoretical scalable, error-correcting Artificial Neural Network

Ciarán Kelly*
Artificial Life Laboratory
RINCE Group
Dublin City University
Glasnevin
Dublin 9
ciaran.kelly9@mail.dcu.ie

ABSTRACT
BrainOS is based on a cut-down version of the Linux operating system. It's main purpose is to facilitate a large, scalable, universal Artificial Neural Network (ANN). The BrainOS handles the interconnection of artificial neurons, and allows for arbitrarily large numbers of these interconnections. It is scalable across multiple machines thanks to its clustering capabilities. BrainOS provides an interface to the internal ANN to any number of inputs and outputs. BrainOS handles all internal aspects of the ANN and is programmed to optimize the interconnections between artificial neurons by reorganizing them during its regular sleep time, somewhat similar to a dreamlike period. BrainOS can be configured in realtime by modifying its various parameters, for example `sleep period`, `neuron distribution` and so on. The goal of BrainOS is to simulate an infant brain, which, through interaction with its environment, can learn something relating to this environment.

1. Introduction
BrainOS is an interface to a large, scalable, universal ANN[1]. It's design is based on a cut down clustering Linux operating system. The BrainOS itself handles all access to the Artificial Neurons held within the system. Using various optimizations, to be discussed later, BrainOS has the ability to re-organize its internal structures to give the Artificial Neurons the quickest possible access to each other. The interconnections between neurons are initially unimportant. It is, however, important to ensure that enough[2] interconnection is initially part of the system. Previous research [1][2] on infant brain development has shown evidence that biological brains begin with an initial "over-fitting" of connections, with, for example, many motor-neurons giving output stimulus to a single motor-actuator, or many sensory nerves feeding into a single neuron. These one-to-many connections later converge to one-to-one relationships as the unneeded connections are pruned, or severed during the early days/weeks/months of the brain's interaction with the environment. These ideas may explain the initially uncoordinated attempts by infants to perform complex tasks such as speaking, walking, writing and drawing. In BrainOS, the unneeded connections are determined and severed through a process of statistical culling. The connections most used are strengthened, while those rarely, or never used are gradually removed. This process is analogous with the concept of "blind variation with selective retention" (BVSR), arguably the driving force behind evolution. While clear for any evolutionist with Darwinian ideals, BVSR has also been postulated to have an effect perception, the immune system, and in this case, the configuration of networks of neurons. Initially, the interconnections are distributed through a process of blind-variation. Subsequently, some of the interconnections are removed through a process of selective-retention.

*Student Number: 99290146
[1]Artificial Neural Network
[2]for the purposes of BrainOS, enough is determined to mean "*as many as possible*"!

2. BrainOS – The Kernel
As mentioned earlier, BrainOS consists of a stripped down Linux kernel. Linux is chosen for it's open nature, allowing us to (reverse) engineer a working Brain kernel. The BrainOS is responsible for mapping the interconnections between the artificial neurons. These connections can occur between neurons on the same physical machine, or can occur across a network. For networking purposes, and to build in some sort of error-prone tendencies, inherently lossy protocols such as UDP are used. UDP is known as a "fire and forget" protocol, in contrast to TCP/IP, which guarantees delivery and acknowledgment of data packets. It is envisaged that BrainOS could "learn" to adapt to such lossy connections by rearranging the neurons that it deems important so that they connect in a more reliable fashion – this could be anything from placing the neurons on the same physical machine to placing them on separate machines that are known to have low communication-error rates. For this purpose, BrainOS can be thought of as a neural router, taking signals from one neuron/sensory-nerve and passing them to another.

The second purpose of BrainOS is to open an interface from the outside environment to the neural network. This interface is simply designed, and completely abstracted from any particular purpose other than getting an input/ output signal to/ from a neuron, or group of neurons. Blind variation again takes place here, with the interface blindly mapping signals to and from neurons, preferably ones that have not stabilized for any specific purpose. An example of such outside world interfacing could be that of both a loudspeaker and a microphone connected to the system. These devices are complementary for the purposes of this experiment as they both deal with sound waves. The BrainOS interface could provide for a discretised signal mapping between the microphone and some neurons, and the speaker and some neurons. It is not hard to imagine that the neurons will be firing in a somewhat random fashion when the microphone is placed in a "noisy" environment. In this situation, the system is becoming aware of the sonic inputs and, in a manner of speaking, becoming conscious of the fact that it can "hear". Of course, the system will not specifically acknowledge this feat as "hearing", it just knows that it is receiving an input and that the input is constantly changing. Now imagine that the loudspeaker neurons are going through "random" firings, not because they are explicitly programmed to do so, but because of the blind-variation in their connections, they are being triggered by outputs from initially unrelated neurons. It can easily be imagined that with such a large interconnection space, some of the firings to the loudspeaker could also be sent to the microphone, to be "perceived" differently. Imagine the situation where the loudspeaker outputs a signal, or "makes a noise" in human terms. If the microphone were to perceive that noise, and subsequently recognise that the system itself was responsible for making that noise, it would now become "aware" of the relationship between the loudspeaker and microphone. It could be described as a sort of positive feedback loop, with increased output leading to increased input, which, afterall, is what the system depends upon. As anyone who has spent any amount of time in a room with a screaming baby will know, there is only so much crying one can take before action is taken to "calm" the situation. Without advocating cruelty to children, the BrainOS system could be administered some form of punishment to inhibit this behaviour, leading indirectly to the system becoming aware of a further input from the environment, brought about by making noise. From the point of view of a baby, this inhibitory attention usually comes from a parent, recognised by the baby as a "giver of things". Communication emerges from such events as the baby learns that certain sounds "cause" certain interactions. The essence of the BrainOS system is to allow the massive interconnection between neural cells that allows such complex behaviour to emerge.

3. BrainOS – Input/Output

It is clear from the example given above that a system such as BrainOS relies heavily on inputs and outputs. Also alluded to above was the need to use discretised inputs and outputs. This is required so as to limit the number of perceptual pathways needed. It could also be argued, however, that merely discretising the inputs could suffice, with the outputs being modified by other neurons in a "fine-tuning" fashion. In the above example, the microphone interface could be set up to recognise discrete inputs on a scale measured in units of 100Hz, covering a specific spectrum of frequencies. The loudspeaker, on the other hand could give a broader range of outputs as the discrete, 100Hz, units are perturbed by signals from other neurons. This provides much food for thought, as it gives rise to the situation whereby the system learns to recognise its own voice. Something that is important to BrainOS, and in fact so important that I will repeat it many times throughout this paper, is the complete disregard by BrainOS to any function other that facilitating the interconnection between neurons and other neurons, and the interconnection between neurons and input/output sensors/actuators. BrainOS serves as a gatekeeper between the external environment and the internal neural network. It has been argued [3] that it is extremely difficult, if not outright impossible to completely analyse or debug any artificial neural network consisting of more than a dozen or so neurons. Even if it were possible to analyse completely a network of a hundred or even a thousand neurons, we are talking here about a massive network, containing as many neurons as is physically possible, given the limitations[3] of modern computers. It may, however, be possible for BrainOS to provide some sort of 3D representation of the configuration of the neurons, perhaps by using a spherical model to represent the total number of neurons, and positioning individual neurons relatively close to the other neurons they interact with. This would give human observers the opportunity to detect emergent structures and maybe even sever connections or entire collections of neurons to observe the results.

To get back to the point of BrainOS input/output, it is not important to have some sort of preconceived notion of what a particular input or output means to us, all that is important is allowing BrainOS to interact with the sensor[4] or actuator[5]. The BVSR[6] capabilities of BrainOS should take care of any internal representation that is needed, or indeed even possible, without any need for human intervention.

4. Speculative Applications

The example usage of BrainOS as layed out above could, theoretically, give rise to language learning and subsequent usage of that language by the system to further its own needs. One only has to look toward human babies, with their initially over-configured brains to see that this kind of BVSR can lead to emergent properties such as language. This would, however, be extremely arrogant on our part, given that the best, if not only, known example of this is in human babies, and this itself has taken billions of years of evolution and at least one year of real world, day to day, minute to minute interaction with other humans who care for the wellbeing of the baby. It would be more appropriate to concentrate on lower level applications such as motor movement and locomotion for our initial explorations. This is not to say that we should not strive to implement the best possible design for BrainOS, with its massively interconnected numbers of neurons. It could be argued that this level of system is essential to producing interesting output. It is also worthwhile to research the best methods for imple-

[3] memory, CPU speed, latency

[4] eg. the Microphone

[5] eg. the Loudspeaker

[6] blind variation with selective retention

menting such a system to save time later, if/when interesting results have been noted and the system is expanded.

As a simpler example than the loudspeaker/microphone configuration discussed earlier, let us now treat of a motor/locomotion example. Let us first assume that there exists some concept of food or nutrition for the system. This need not be explicitly programmed into the BrainOS, as by its very nature, BrainOS will "decide" what it wants or doesn't want. This food or nutrition concept could be thought of as energy. When there is lots of energy in the system, there is lots of neural interaction, which, at the most fundamental level, is what BrainOS "likes". In the initial stages of development, an external entity, let us assume a human "parent" in this example, brings energy to the system in the form of an electrical cable with a connection jack on one end that connects in some way to the system. The system also has a set of motors which control wheels/limbs, but the system is initially unaware of any correspondence between these motors and the concept of locomotion. Now, let us assume that the energy source also carries a transmitter. This transmitter gives out a constant signal which BrainOS perceives as an increasing signal when the energy source moves closer, and decreasing as it moves away. In this situation, BrainOS should theoretically "learn" to recognise when it is about to receive some energy, due to the firings of some neuron or group of neurons connected up to the signal receiver. After some period of time, the "parent" decreases the frequency of these energy-giving visits to the system[7]. At the same time, when interesting patterns of motor movement are noted, the parent administers an extra energy-giving visit. The system could then learn that by moving limbs, it is possible to "cause" itself to get more energy. At some future stage, it could be imagined that the system could learn to go get the energy for itself, becoming independent of the "parent".

Although this is a simple example, with the energy source appearing to emit an exponentially increasing signal as it approaches and decreasing when it moves away, it might be interesting to consider what might happen if the system used a camera, as an eye, to detect the energy source. Before discussing that possibility, let us first look a realworld example [3] of a neural net in action. In the 1980's, the Pentagon commissioned a project to build a neural network to analyse images to detect the presence of tanks hiding in trees. Two hundred photographs were taken, 100 of trees with tanks hidden in them, and 100 of just plain old trees. Half of each of these groups were retained as a control group, while the other half were used to train the neural net. After some initial random behaviour, the neural net learned to recognise those pictures of trees with tanks in them and distinguish from those without. Then, the control group of photos was retrieved and presented to the system, whereupon it carried on with it's correct detection of the tanks in the trees. The Pentagon, skeptical of such a "perfect" system, ordered another 200 photographs to be taken and presented to the system. This time, instead of correctly recognising the photos containing tanks, the system once again "descended" into randomness. After some brain-bashing and attempts to debug the system, it was "discovered" that in the initial set of photographs, the ones with the tanks in the trees were taken on a cloudy day, while those without were taken on a sunny day. The system, therefore had learned to recognise a cloudy sky, rather than the tanks hidden in the trees! Chalk one up for nature! Getting back to the idea of using an artificial eye for the system, it could be easily assumed that the system would begin to recognise that it was the parent who brought the energy, and respond only to the parent. This has obvious parallels with the natural world when you imagine a nest full of hungry chicks with their mouths opened wide on the return of the parent. As they mature though, they begin to recognise the food as a separate entity to the parent and learn how to get it for themselves. It would be interesting to see if a system such as BrainOS could facilitate such a learning process.

[7]it would be interesting to note whether these visits need to be regular or not

5. Conclusion

In the proverbial sense of things, I have jumped to this conclusion. The reason being that I didn't want to clutter up any clarity that might exist in the paper with any more unfounded hypotheses. Also, there have been many oversights on the part of the examples I have given, such as the lack of a balancing mechanism for the locomotion system mentioned earlier. The human brain is universally accepted to be a highly complex organism. Following many years of research and examination, we have only been able to scratch the surface of its capabilities. It is my opinion that we have, in some ways, opened this Pandoras Box just wide enough to catch a glimpse at whats inside, but lack the technology or knowledge to open it any further. I also feel that we are in an excellent position to attempt to design our own box, from what we know of the original, and I fully expect it to be as complex an organism as its biological forerunner. It is perhaps a huge leap of faith, to accept that something designed by our own hand, could subsequently puzzle us as to its operation, but this very alien concept has shown itself time and again through research in all but the most basic of neural networks. [3]

6. REFERENCES

[1] Jean-Pierre Changeux. 1985, p. 248.

[2] Gary Cziko. Without miracles: http://faculty.ed.uiuc.edu/g-cziko/wm/.

[3] Neil Fraser. September 1998: http://neil.fraser.name/writing/tank/.