Native Architecture for Artificial Intelligence

Dushan Balisson, Wim J. C. Melis, Simeon Keates

Dushan.Balisson@greenwich.ac.uk, Wim.J.C.Melis@greenwich.ac.uk, S.Keates@greenwich.ac.uk

Department of Engineering Science, University of Greenwich at Medway,
Chatham Maritime, Kent ME4 4TB, UK

Introduction

The brain is a complex organ and even to this date, very little is known about how it works. Through the years, replicating the intelligence of the brain has puzzled many scientists, and most of this work can be broadly classified into two categories – Neurophysiological or Cognitive based. The latter approach tends to overlook the actual structure of the brain in order to focus on the behaviour itself. This is e.g. exemplified by the Turing test [1], which implies that if a human interacts with an artificial machine and identifies it as a human, then this artificial machine is considered sufficiently humanlike.

Proponents of the neurophysiological approach argue that the intelligence of the brain lies in its structure, hence if this structure can be replicated, one should be able to replicate human intelligence. In this context, one of the most complete neurophysiological models of the neuron is the Hodgkin-Huxley model [2], which has served as a reference for biological plausibility of subsequent neural models. However, a major issue with the Hodgkin-Huxley model lies in the complexity to use it for the implementation of a complete and useful network. On the other extreme, one can find simple models such as the Leaky Integrate and Fire (LIF) model [3], which is the most widely used neural model, but is commonly regarded to be oversimplified. While this model is biologically implausible, it is computationally viable and can therefore be implemented into relatively large networks to study their behaviour and dynamics.

Even further simplified networks are already used to feed the need for intelligent machines that are able to learn. Today, one of the most successful machine learning paradigms is the Artificial Neural Network (ANN). ANNs implement a self-improving function by giving a certain output, based on certain inputs, where each neuron is modelled as a transfer function. These systems have shown commendable performance and form the basis for certain popular services, such as the Google Brain Project that provides YouTube users with recommended videos based on their viewing history [4]. However, the computational complexity of these tasks is not to be underestimated, as the Google Brain combines 16000 computers to deliver the capabilities of a rat's brain [5]. This raises the question as to why our current systems are so largely inefficient when it comes to delivering brain-like functionality. Considering that the theoretical limit of silicon feature sizes is getting closer and that the collapse of Moore's law is upon us, it seems more than necessary to get back to the drawing board and start considering alternative platforms that are more efficient and run without the huge sematic gap.

Method

Since the initial conceptualisation of the ANN framework, much more is now known about cognition and neurophysiology. While this knowledge covers various levels, the aim here is to look at the brain from a higher abstraction level and use that to identify the main function of the brain. This approach can also salvage knowledge from our understanding of the ANN model. Therefore, as an example, an abstract model of the auditory pathway, as shown in Figure 1, will be investigated in more detail.

The inner ear, more specifically the cochlea, transduces the received frequencies into vibrations of frequency selective hair cells; these produce electrical impulses that are then fed to the brain. The nature of this signal can be extrapolated from what is known about the functioning of the auditory pathway, however the actual representation of the information is purposefully overlooked here to maintain an abstract perspective. Within the brain, the processing that takes place is closely linked with the nature of the input signal, and so as a first step it is essential to understand and break down the processing that takes place in the brain. Being an associative memory system, the brain compares the incoming information with what is known, through similarity matching. Similarity matching is one of the most important tasks in the brain and the brain seems to excel at this task.

Figure 1: Abstract Level View of the Auditory Pathway

Implementing similarity matching can obviously be achieved through Euclidean distance calculation or various other mathematical/statistical approaches, however these seem rather brute force approaches and may not be the most efficient in their own right. Additionally, one may need to consider how information can be stored and processed within the same location to allow suitable encoded sensory information to interact with the stored information. Obviously, storing every single possibility of a sensory item is implausible because of its impact on the required amount of memory, so one solution could be to use a hierarchical structure. This would mean that the brain breaks down its comparison into stages with different abstractions [6] and so comparison at one level is performed before comparing at the next level which has a higher abstraction, and so on and so forth.

To make such an artificial system as efficient as possible, there is the need for sensory information to interact with all possibilities simultaneously at a particular hierarchical level, which then again brings along several challenges. At the same time, it is expected that this approach needs to allow learning in its broadest sense. Therefore, the memory and processing structure should be modifiable through a learning and feedback mechanism in a way similar to, but not limited by, a backpropagation algorithm.

Discussion

Today conventional computing platforms are inherently focused on arithmetic and logic operations, which seem a long way from what happens within the human brain. While there seem to be a variety of different similarity matching based approaches around, having to perform them on a one by one comparison basis, either sequentially and/or with a certain amount of parallelism has performance implications. It therefore seems essential that effort is put into the development of platforms that model the brain from a functional perspective and provide for a more direct mapping to technology. That being said, until other platforms become a reality, conventional computing platforms will remain an indispensable tool to support the exploration for alternative artificially intelligent platforms.

References

[1]   M. Alan, "Turing. Computing machinery and intelligence," *Mind*, vol. 59, no. 236, pp. 433–460, 1950.

[2]   A. L. Hodgkin and A. F. Huxley, "A quantitative description of membrane current and its application to conduction and excitation in nerve," *Bull. Math. Biol.*, vol. 52, no. 1–2, pp. 25–71, Jan. 1990.

[3]   E. Orhan, "The Leaky Integrate-and-Fire Neuron Model," no. 3, pp. 1–6, 2012.

[4]   P. Covington, J. Adams, and E. Sargin, "Deep Neural Networks for YouTube Recommendations," in *Proceedings of the 10th ACM Conference on Recommender Systems - RecSys '16*, 2016, pp. 191–198.

[5]   Q. V Le, M. Ranzato, R. Monga, M. Devin, K. Chen, G. S. Corrado, J. Dean, and A. Y. Ng, "Building high-level features using large scale unsupervised learning," *Int. Conf. Mach. Learn.*, p. 38115, 2011.

[6]   J. Hawkins and S. Blakeslee, "On intelligence," *Neural Networks*, p. 272, 2004.