

Biologically Plausible Hybrid Network Design and Motor Control

Gregory R. Mulhauser

Department of Philosophy, University of Edinburgh
Edinburgh, Scotland EH8 9JX
scarab@ed.ac.uk

Abstract. I outline a general theoretical framework for understanding the relationship between genetic algorithms and traditional learning methods on the one hand, and between symbolic and distributed information processing on the other. By way of example, I discuss aspects of motor control problems and describe a theoretical biologically plausible hybrid network inspired by the theoretical framework. Preliminary analysis suggests the net might be a useful point of departure for developing biologically plausible visually guided motor control systems.

1. Introduction

Present development of artificial neural networks, whether by genetic algorithms or by traditional learning methods, is grossly inadequate as a picture of biological reality. Network design with genetic algorithms ignores the rôle of ontogenetic behaviour adaptation in response to the characteristics of the particular environment in which an individual phenotype finds itself. Genetic algorithms typically are applied to generate zero plasticity networks incapable of ontogenetic development. (This is the case even for so-called 'growing networks' such as [1] in which the ontogeny of a network is directly, albeit nonlinearly [2], encoded in the genotype and is independent of environmental factors peculiar to particular individuals.) Conversely, network development with traditional learning methods generally ignores the fact that real biological organisms learn within boundary conditions set by the organism's genotype and phylogenetically adapted by recombination, mutation, sexual selection, and environmental pressures. The boundary conditions of traditional learning networks are set not by the power of evolution but by human designers taking 'educated guesses' at appropriate architectures, learning algorithms, and connection patterns.

But it scarcely needs pointing out that real adult phenotypes are the product of both ontogenetic *and* phylogenetic development. The living organisms we encounter every day are genetically endowed at birth with a wealth of characteristics which evolution has determined are beneficial for their survival and reproduction, but within the bounds laid down by their genotypes they are also capable of adapting to changes in their environment. They are capable of *learning*. Thus, insofar as the theoretical frameworks offered by either the genetic algorithms paradigm or the traditional learning paradigm claim to be biologically plausible, they are incomplete, and insofar as they claim to be complete, they are not biologically plausible.

2. An Alternative Framework

I suggest a broader theoretical framework within which artificial neural network design mimics the natural features of both genetic coding and environmentally induced adaptation by individuals. On this view, genetic algorithms should be applied to genotypes which code not only for the standard parameters describing nodes, thresholds, and connections (or simple dendritic growth for growing networks), but also for characteristics of dendritic plasticity. These characteristics might include the learning rate of a globally defined Hebbian [3] rule as well as a measure of the capacity for growing new connections. The fitness function may then be applied to a *phenotype* grown nontrivially during a 'learning phase' in a dynamic environment. It might give preference to phenotypes which adapt smoothly to environmental changes, such as a motor control network which could not only manoeuvre a robot arm toward a given point but could also adapt to avoid obstacles introduced into its path.

From artificial neural nets developed within this theoretical framework, I believe we may gain some insight into the emergence of classical symbolic processing in real intelligent organisms. I propose that evolution creates hybrid architectures of high and low plasticity connections in which arrays of neurons with low plasticity connections might embed very simple classical processors such as basic logical connectives. This view takes theoretical support from the idea that low plasticity subnetworks implementing classical functions may be the most efficient 'building blocks' on which higher plasticity distributed hybrid networks could rely. Given that genotypes specify at least some, if not all, characteristics of dendritic growth (by virtue of specifying the structures of cells themselves), I believe it is highly *implausible* that Nature could have failed to exploit this elegant way of mixing the best attributes of distributed and (neurally implemented) classical systems.

This is not to say that genetic algorithms may *only* generate networks exploiting classical processes, for this is clearly not the case. If adult organisms never had to learn by experience, never had to remember information or respond to situations radically different from those which influenced the phylogenetic development of their predecessors, genetic algorithms might have provided for the entire repertoire of behaviour of adult organisms with nonlearning networks operating with any balance of obviously symbolic or distributed principles. But in the real world higher organisms are not entirely hardwired by their genotypes. I suggest simply that real biological development yields networks of mixed plasticity and that in those portions of phenotypes which *are* hardwired, we may find embedded symbolic functions.

By way of example of what might be accomplished by deliberately mixing high and low plasticity connections, in what follows I describe a motor control problem together with a manually designed speculative prototype network meant to illustrate some architectural principles which might be exploited by the automated development theory I have described above. It is certainly inferior to what could be generated by such a strategy, so I include it not as a solution to motor control problems but merely as an example of a first step.

3. The Sample Problem

I am concerned with the problem of visually guiding some mechanical device to an arbitrary point in space. This amounts to combining information about the present visual image with information about the desired image in order to activate a motor

system. Here we assume that the relationship between a given activation of the motor system and its influence on the visual image is initially unknown but could be described by a (hopefully simple) computable function. For this example, we also assume that visual information has been pre-processed in such a way that the control system is presented with an indication of, for instance, the present real coordinates of the device together with its desired coordinates. The dimensionality of the problem might be increased by also including the present and desired coordinates of more than one coupled part of the mechanical device, such as both the target end and the elbow joint. The details of the specific device under motor control do not concern us. Instead, I would like to paint in broad strokes a picture of one possible neural architecture for performing this type of control task.

3.1 Architecture

The network I propose works on the hypothesis that an unsupervised Hebbian network provided with feedback about its level of success at performing a particular task might approximate the capabilities of a supervised network learning to perform a similar task. This principle is inspired by Rumelhart's biologically plausible implementations of something similar to backpropagation and his recent use of this type of network in motor control and so-called 'mental mapping'. [4] Rumelhart's networks are primarily Hebbian but rely upon feedback from nodes producing a theoretical neuromodulator which regulates plasticity without affecting activation. The neuromodulator is plausible, but it has yet to be identified in biological systems; the present network is meant to perform a similar task without recourse to this modulator.

Rather than providing a plasticity-modifying chemical at a particular neural junction, the present strategy is simply to provide additional excitatory input to the two Hebbian nodes in question. In most networks using correlation rules to update weights, such as that described in [5], high correlation between node outputs must be combined with high present output frequency to achieve a maximal update to the connection strength. Thus the strategy of providing additional excitatory input to the two nodes increases the magnitude of the connection update. Of course this will also influence the other nodes to which either of the two in question might be connected, so the strategy is far from identical to the neuromodulator scheme.

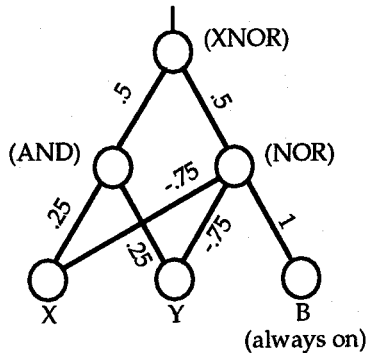


Fig. 1. XNOR Function

The network receives its feedback about performance from a zero plasticity sub-network which provides a classical measure of the improvement in position caused by the most recent motor activation. A rudimentary version of such a measure is the XNOR function, implemented as shown in Figure 1 by perceptron-style units with all or nothing thresholds set to .5.

However, since we require more information than a simple verification of whether two units are either both on or both off, a more flexible measure is the function $1 - |X - Y|$, implemented as shown in Figure 2, together with its output graph. Here the nodes have zero thresholds and a continuous output response which can be read off from the cross section where the graph meets either of the two vertical planes made by the axes. I will refer to this simple network as a 'convergence detector'; it may be nonlinearised by altering the output function of the final node to match the cross section of the desired graph.




Fig. 2. Convergence Detector and Output Graph

In the complete system, depicted in Figure 3, convergence between the present image and the desired image is measured by the detector A and convergence between a previous image (thus the propagation delay) and the desired image by the detector B. While each of these detectors has a number of output signals identical to the number of dimensions of the image information, for simplicity only one output each is shown here. The outputs of the two detectors are compared by single nodes with smooth output functions and complete efferent connections to the control network which effectively subtract the old convergence from the new in the case of the subtractor marked C or the new convergence from the old in the case of D. The output functions of these nodes must be scaled to amplify positive results: since we would expect sudden very high improvements in convergence only rarely, it is important to magnify the presence of even minor improvements to a significant level. The D node is unique in that it represents a neuron which produces only inhibitory chemicals, but its efferent connections are still updated according to the Hebbian rule in place. (Note that some sort of arrangement with fixed afferent connections and plastic efferent connections is necessary for communication between nodes with fixed connections and nodes which are part of a learning network.)

The upshot of this system is that when the controlling network yields a motor output which improves the convergence between the desired and the actual visual input, all the nodes in the control network will receive additional excitatory input from C, but only the connections with those which were firing will be strengthened.

When the motor output worsens the convergence, all the nodes receive inhibitory input from D, but again only the connections with those which were firing will be strengthened.




Fig. 3. Complete System

Additional excitatory input to a pair of firing nodes which are connected to each other in the main network will contribute to an increase in the strength of the connection between them, whereas inhibitory input will contribute to a decrease in the connection efficacy (or at least a relatively smaller increase). The interaction between the inhibitory and excitatory inputs from the subtractors themselves and the influence of this input on the connections between the nodes within the main network is meant to be a primary mechanism affecting the system's development.

3.2 Comments and Shortcomings of the Architecture

Perhaps the most telling criticism against the system as it stands is that there is no guarantee the main network will converge when it is trained upon a set of desired images with a consistent environment and consistent motor actuator characteristics. As I have stated, however, the architecture is intended primarily as an example of the kinds of features which might emerge from the design theory I described above, and it illustrates the application of a hybrid architecture to the feedback hypothesis. A more subtle criticism is that complete efferent connections from the subtractors suggest that in terms of reinforcement no discrimination will be made between nodes whose firing was highly desirable for achieving a convergence improvement and those whose firing was only marginally desirable; all nodes whose firing contributed to a positive change in convergence are reinforced, while all those whose firing contributed to a negative change are inhibited. Improvement of the network to remedy these deficiencies and others awaits full implementation, testing, and detailed analysis.

4. Discussion

I have outlined what I believe is the most powerful theoretical framework for biologically plausible neural network design considered to date. I have given reasons for broadening the class of neural network characteristics under the control of a genetic algorithm as well as reasons for incorporating a 'learning phase' of ontogenetic development and for applying the fitness function to the phenotypes thus produced. I have indicated how this theoretical framework may allow us to view some phenotypes as hybrid networks in which evolution has 'discovered' useful classical functions and embedded them in low plasticity subnetworks. I have given an example of a problem which could benefit from this type of approach and described a manually designed sample network which might be useful for solving the problem.

For complex tasks, the capabilities of networks produced under this theoretical framework may surpass those of most artificial networks either genetically created or manually designed for ordinary learning. Since artificial networks are not constrained by the boundary conditions of biological neurons operating in real space, such artificial networks may ultimately surpass the capabilities of similarly connected biological counterparts. Given computationally simple correlational learning algorithms and ontogenetic phases of a few thousand cycles or less, this type of neural network design is well within the bounds of existing technology.

Already the networks which emerge from simpler genetic algorithms are extremely difficult to analyse in terms of the functions of individual neurons. The architectural subtleties made possible by the framework I have described may prove still more resistant to functional analysis. This makes yet more pressing the need for ongoing consideration of the problems created by automated generation of more and more capable neural networks which we understand less and less.

References

1. Stefano Nolfi and Domenico Parisi: Growing neural networks. Presented at Artificial Life III. Santa Fe, New Mexico: June 1992.
2. C.G. Langton: Artificial life. In: L. Nadel and D. Stein [eds.]: 1991 Lectures in complex systems. Reading, Massachusetts: Addison-Wesley 1992.
2. D.O. Hebb: The organization of behavior. New York: Wiley 1949
3. David E. Rumelhart: Unpublished presentation at Fifth JCI Summer School in Cognitive Science and Human Computer Interaction. University of Edinburgh, Scotland: September 1993.
4. Gregory R. Mulhauser: Population coding in a theoretical biologically plausible network. Proceedings of the European Symposium on Artificial Neural Networks 1993, pp. 65-70.

Acknowledgement

I am grateful to HM Government's Marshall Aid Commemoration Commission for financial support of my research during the time this paper was conceived.