Design and Use of AUTONET, a Logic-Level Neural Simulator

by

R. E. Ellis

A thesis
presented to the University of Manitoba
in partial fulfillment of the
requirements for the degree of
Master of Science
in
The Department of Computer Science

Winnipeg, Manitoba, 1981

DESIGN AND USE OF AUTONET,

A LOGIC-LEVEL NEURAL SIMULATOR

BY

R.E. ELLIS

A thesis submitted to the Faculty of Graduate Studies of
the University of Manitoba in partial fulfillment of the requirements
of the degree of

MASTER OF SCIENCE

I hereby declare that I am the sole author of this thesis.

I authorize the University of Manitoba to lend this thesis to other institutions or individuals for the purpose of scholarly research.

R. E. Ellis

I further authorize the University of Manitoba to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

R. E. Ellis

The University of Manitoba requires the signatures of all persons using or photocopying this thesis. Please sign below, and give address and date.

ABSTRACT

AUTONET is a system for the specification and simulation of neural networks. It deals with the network as a directed graph, the nodes of which are automata which accept a tree of inputs and produce a single output, and the lines of which are queues of arbitrary length via which the cellular automata communicate.

The simulation system includes the WARP network construction language. The important features of WARP are its ability to manipulate not only single cells, but also arbitrary groups of cells, as distinct units. Projection of one cell group upon another is facilitated by this feature, and also by the ability to permute the interconnections between cells. This permits convergent and divergent projections in many dimensions to be handled with ease.

The experiments conducted with AUTONET have been largely with simple and moderately complex networks. Some apparently trivial networks are shown to actually possess a wide variety of possible rhythmic outputs. These output patterns are characterized as to their stability, and to the connectivity patterns which make the various patterns more or less likely to be generated in a network subject to random effects.

Considerable simulation of the gastric and pyloric subsystems of the lobster stomatogastric ganglion was also done. It is shown that all of the rhythms generated by the ganglion may be explained by decomposing the networks into groups of two or three cells, which smaller networks have already been well-characterized.

Evidence is also presented for the hypothesis that the endogenous bursting cells of the pyloric subsystem are not the generators of the pyloric rhythm. Rather, their purpose seems to be as synchronizers of the cycle.

ACKNOWLEDGEMENTS

I wish to greatly acknowledge the assistance I received from my friend and adviser, Douglas Young.

Thanks also are due to Jon Muzio for both syntactic and semantic observations, and to Patty Churchland for the many stimulating conversations which led to my interest in lobster stomachs.

This thesis was made possible in part by a scholarship from the Natural Sciences and Engineering Research Council of Canada.

Chapter I

INTRODUCTION

Among the many lesser-researched areas in artificial intelligence is the subject of simulating neuronal networks. Just as man has surpassed his body's abilities in other ways, the general feeling goes, so he might be able to do so in the processing of information. A crane does not operate exactly as the arm does, yet lifts greater weights and works for days on end; perhaps by discovering the fundamental principles of thinking (as he discovered the principles of the six simple machines) man might yet construct a device that is truly intelligent.

Viewed in this light, neural simulation seems a curiosity. It seeks to know how things <u>are</u>, instead of how they <u>ought</u> to be. This research paradigm is certainly of concern to biologists, and perhaps even to psychiatrists, but since the modern computer is quite unlike the brain, in either structure or operating principles (such as the brain's are known), there seems to be little direct relation between them.

The rejoinder to this argument is that determining the operation of a network of elements individually as complex as neurons is an extraordinarily difficult procedure. It is

a problem in the analysis of information flow, which the computer scientist is uniquely qualified to study. It is so complicated that pure analysis is largely impossible because of the paucity of accepted first principles, and only simulations with computers (especially large computers) have regularly yielded consistent, interesting results. And neural networks are highly parallel, with many feedbacks; since most formal theory of computation deals with sequential machines, any principles of operation for networks would constitute a contribution to a little-explored but increasingly active field of research in theoretical computer science.

Also, the main work to date has been done by biologists interested primarily in explaining the detailed operation of selected networks (although they of course hope for a grander understanding, they recognize that it is not immediately forthcoming and hence concentrate on whatever they can gather the most data about). There is not even a catalogue of the possible activities of simple networks of commonly-investigated cells. The main impediment here is, once again, the enormous complexity of even the smaller nets.

The AUTONET simulation system is an attempt to make such basic research more approachable. By way of introduction, it is necessary to examine the structure and function of neurons; only then does one understand why so many resources

have already been expended on neural simulation the world over. At one time or another, all of this introductory material has been incorporated into someone's simulator. This zest for all-explanatory power can produce such a plethora of conditions, and such a wealth of data, that many simple patterns were for some time effectively obscured. This neurological data is not presented solely to further an understanding of the AUTONET system (although the terminology is useful); it is presented more so that the reader can better grasp why the high level of analysis represented by the AUTONET approach has been adopted, i.e. so that one can see just how much data there is to obfuscate the research.

After the neurology will come a chapter on AUTONET itself. The primary problem in analysing networks (and thus primary expenditure of energy in designing AUTONET) has been that there is no convenient, widely-accepted language in which to express the network. The WARP network construction language thus constitutes an examination of some of the needs of high-level simulation languages, and suggests some ways of achieving these needs. There are perpetual tensions amongst the competing ideals of completeness, conciseness, clarity, convenience, and computational efficiency. How these are resolved involves techniques drawn from the disciplines of formal languages, compiler design, information theory, and man-machine engineering; there is no

common language shared by all of these, so some terminology is also presented.

The system described, one then moves to a chapter which discusses networks *per se*. This is a brief study of how networks and their component cells might be designed, and of the interactions present in even apparently trivial networks. Following this is a thorough simulation of the network found in the stomatogastric ganglion of the lobster. This network has been much studied elsewhere, but even it yields a few unexpected results.

Chapter IV also displays the value of experimental evidence, gathered via extensive computer simulation, in the analysis of biological systems. In particular, it is possible in a simulation to rapidly modify the network parameters, leading the investigator to a better comprehension of the network's operation.

In such simulations, it is traditional to implement some sort of adaptivity in the network elements. This issue is carefully considered, and the conclusion drawn is that given the enormous complexity of apparently simple systems, a better understanding of the operational principles at work is a prerequisite to proper adaptive algorithms. The emphasis is thus more on the operation of the robust peripheral nervous system than the intricate, adaptive central system.

The thesis concludes with some closing remarks on network design, and what might be expected from future simulations.

Chapter II

INTRODUCTORY NEUROLOGY

The neuron is a eukaryotic cell, found in many organisms, that is characterized chiefly by the fact that its transmembrane potential can change by tens of millivolts in milliseconds due to the presence of pores in the cellular membrane, which permit the passive flow of the ions responsible for the transmembrane potential in the first place. The potential change can be facilitated by another neuron releasing chemicals into a synapse; by a close abutment of neurons electrotonically inducing a potential change; or by the neurons breaching their cell walls to form virtually continuous membrane. A receptor cell, which transduces (and usually amplifies) energy such as light or sound, can also have selective effects. No other natural selective effects are known, although there is considerable research into non-selective effects such as hormones and drugs.

To speak of <u>the</u> neuron is somewhat misleading, for there are at least several dozen types found in primates alone. The number of types in more specific, less plastic nervous systems, e.g. the invertebrates, is truly staggering. Thus, when discussing neurons, unless otherwise specified it may

be assumed that one is dealing with the primate central nervous system (CNS) pyramidal cell, which is the primary constituent of the cerebral cortex.

2.1 MORPHOLOGY.

Neurons[1] vary in the size of their somata from 2 to 1000 microns, pyramidals from 15 to 40 microns. They are notable for having processes, extending from the soma, which are at least several times the somatic diameter (the cell innervating the adult human big toe has a process over one metre long).

The nucleus of a neuron, like that of most secretory cells,[2] is much larger than is typical in multicellular organisms, often consuming half the somatic volume (79% of a linear dimension). As is usual in eukaryotes, the nucleus is enveloped in endoplasmic reticulum, and most of the genome is contained in the nucleolus. However, neurons are unique in that it is exceedingly rare, in higher species, for an adult cell to reproduce; concomitant with this absence of mitosis is a thin and wispy chromatin, as opposed to the robustly-staining variety normally encountered.

---

[1] This and other general neurology may be found in the excellent introduction by [Shepherd,1974]. More detailed information is available in THE NEUROSCIENCES Study Program series by Rockefeller University Press.

[2] Most neurons excrete transmitter substances, in addition to having trans-membrane ion flow. These two activities require an especially active cellular support system.

Also present in the soma are the standard eukaryotic organelles: the rough endoplasmic reticulum, speckled with ribosomes that employ m-RNA in protein synthesis; the Golgi apparatus, which appears to be active in the "packaging" of substances into saccules, e.g. transmitter substances into vesicles and enzymes into lysosomes which lyse cellular wastes; the microtubules and microfibrils which, aside from constituting the structural ground substance, also actively transport material within the cell; a centriole, known to play some part in organising the microtubule network (especially during mitosis); and mitochondria, ubiquitous in eukaryotes as the seat of cellular respiration, and particularly plentiful in neurons, which consume a lot of energy in signalling.

The cellular membrane, 7.5 to 8.0 nanometres thick, is in most parts a phospholipid bilayer largely impervious to both polar and non-polar molecules. In addition to the normal transmembrane transport associated with staying alive, e.g. glucose in, carbon dioxide out, there are channels of .4 to .8 nm I.D. that allow small ionic species such as hydrated Na^+ and K^+ to pass through the membrane when the normally-closed channels open up. The channel density is, per square micron of membrane surface, 100-200 on the soma and its proximal processes, and up to 800 on parts of the axon. The maximal flow rate of the Na^+/K^+ channels seems to be 200-220 Na^+/sec., 120-130 K^+/sec. per channel.

The cell attempts to segregate various ionic species with respect to the cellular membrane. Because of the ionic charges and relative densities, this produces a potential difference between the intra-cellular fluid (ICF) and the inter-stitial fluid (ISF) surrounding the cell. By examining the potentials and respective concentrations experimentally, and comparing this with the Nernst equations for generic transmembrane potential:

$$E = RT/zF * \ln(O/I)$$

where
- E : potential difference between outside and inside
- R : Rankine gas constant
- T : absolute temperature of system
- z : valence of ionic species in question
- F : Faraday constant
- O : outside concentration of ionic species in question
- I : inside concentration of ionic species in question

then one finds that for the mammalian skeletal muscle cell:[3] which has an ICF resting potential of -90 mV, that the potential if only K^+ ions where present and so segregated is -97 mv (both values taken at 37° C.). It is thus considered sufficient to examine only K^+ flux for most cases.[4]

---

[3] Which is quite similar to neurons, but much larger and hence easier to study.

[4] One can readily calculate from the above, assuming perfect membrane capacitance, that hyperpolarization to +40 mV causes a loss of 1 in 10^7 ions if, as in large axons, each cm^2 of membrane encloses 10^5 moles of K^+. The loss goes up to 1 in 3000 for small neurons.

TABLE 1

Ionic Concentrations in micromoles/ml

SPECIES	ICF	ISF
K+	155	4
Na+	12	145
Cl-	4	120
assorted anions	155	34
assorted cations	1	5

As previously mentioned, neurons have processes extending from the soma. When stained, the dendrites exhibit so striking a resemblance to a heavily-branched tree that it is common in the literature to refer to trunks, branches, and arborization patterns. The axon, by comparison, appears to proceed axially from the soma for quite some distance before arborizing (hence the name).

The morphology of the process membrane is virtually identical with that of the soma. The exception is that the axon, if it proceeds more than a millimetre or two from the soma, is <u>myelinated</u> (dendrites rarely, if ever, are). This myelin sheath is composed of the processes of satellite cells (Schwann cells in the PNS, oligodendrocytes in the CNS). There are few, if any, ionic channels under the myelin; they are at the maximum observed density (800/sq. micron) at the nodes. This myelination makes possible a

fast, sure conductance which will be described functionally below. It is unfortunate, especially for physiologists, that only large axons are myelinated, for distinguishing fine unmyelinated axons from fine dendrites in a virtually transparent three-dimensional agglomeration has proven near impossible, and accounts in part for the paucity of data which may be employed in simulations.

The only other morphological peculiarity of note is that as the axon proceeds from the perikaryon (or from a large dendrite) there is often found a cone-shaped region called the axon hillock; the cone seems to be a funnelling of microtubules and neurofilaments into the energy-consuming axon, and may indeed play a largely metabolic support role. Between the hillock and myelin start is the initial segment, characterized by a dense undercoating of the plasma membrane.

There is some controversy surrounding this initial segment. Not only is it definitely not physically present (in an observable form) in the Purkinje cells of the cerebellar cortex, but it is far from clear that it is present in the small, unmyelinated pyramidal cells which abound in the cerebral cortex (although it is found in many large pyramidals). It seems to be becoming more fashionable to speak of the spike initiating zone (SIZ) of the axon. This is in many ways to be preferred, especially when dealing with the motorneurons of some invertebrates, which neurons may have several SIZ's.[5]