The Robot Builder

Robotics Society of Southern California

Volume Eleven Number Seven July 1999

Notices

- Introductory Mobile Robotics Class - 10:00AM - 12:00PM
- Business Meeting - 12:30 - 1:00
- General Meeting - 1:00 - 3:00

Distribution

If you would like to receive The Robot Builder via e-mail, contact the editor at:

apendragn@earthlink.net

Inside this Issue

Introduction to Neural Nets
Part 1

by Arthur Ed LeBouthillier

Neural Nets excite the imagination of many as an alternative way of creating smart processing systems for use in robots. They offer the possibility of creating intelligent systems without all of the requirements of traditional programming plus they offer the added benefit of being able to be run simultaneously on multiple computers, thus better utilizing parallel processing.

The following is a brief introduction to neural nets. First, we will examine a few of the characteristics of animal neurons, then we will examine a model of them. In a future article, we will review a learning algorithm which can be applied to this model.

Animal Nervous Systems

Animal brains are composed of numerous specialized processing cells called neurons plus other support cells such as glia. The glia are thought to aid in the distribution of nutrients and space fillers between neurons. The neurons provide the processing elements which are interconnected to provide paths for signal flow. There are many different types of neurons which perform varying functions when combined in circuits.

Figure 1 provides a rough overview of the structure of a neuron. For the most part, a signal flows from dendrites through the soma to the axon. In real life, there are potential fields which can be created between closely spaced cells which also affect the cell's firing.

A neuron generally exhibits an all-or-nothing behavior whereby the cell is either at a resting state or else in a pulsed firing state. For some neurons, there is a steady firing rate which is dependent upon the sum of the inputs; other neurons don't fire unless a certain level has been attained on the inputs. Another characteristic which real neurons have is a certain amount of propogation delay between the input and output. Because

(see Neural Net on page 5)

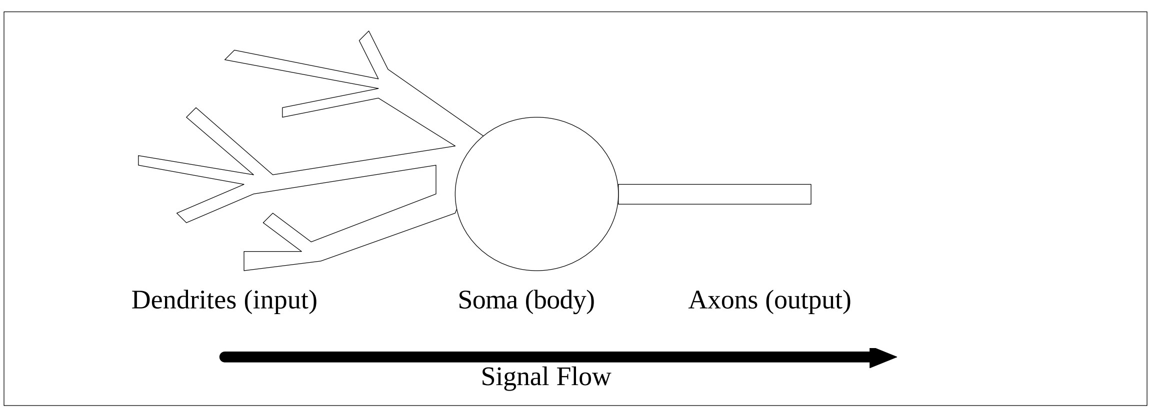

Figure 1 - the parts of a basic neuron

Japanese introduce ant-size robot

(AP) 21 June 1999 - Japanese electronics companies have developed a micro-machine the size of an ant that can crawl around thin pipes, inspect and even fix problems at power plants, officials said Monday.

The box-shaped robot is only 5 millimeteres (0.2 inch) long, 9 millimeters (0.36 inch) wide and 6.5 millimeters (0.26 inch) high. It has a pair of round connectors on both sides that can be linked up with other robots for more extensive assignments. With a weight of only 0.42 grams (0.0147 ounces), the robot can lift objects twice as heavy as itself and can move at a speed of 2 millimeters (0.1 inch) per second, said Koji Hirose, spokesman for the Ministry of International Trade and Industry.

Mitsubishi Electric Corp., Sumitomo Electric Industries Ltd. And Matsushita Research Institute Tokyo, Inc. developed the machine under the government's 25 billion yen ($206 million) "micro-machine" project, Which began in 1989, Hirose said. The robots are one of three types of machines designed for use in different environments, he said.

The robots, which can crawl into the tiniest gaps around bundles of pipes, are expected to speed up inspection and repairs at electric and nuclear power plants because they can be sent in while the plants keep running. Scientists are working to add new functions to them so the robots can climb up and down a pipe while connected to other machines. They also plan to develop robots with motors and problem-detecting sensors.

New Mars Rover tested at JPL

A new Mars rover is being tested at JPL. Named FIDO, for Field Integrated Design and Operations, it is a prototype of the Athena rover planned for the Mars sample return mission scheduled for 2003 and 2005.

The new rover is substantially larger than the Sojourner rover of the Pathfinder mission. Weighing in at over 150 pounds, this prototype is nearly twice as large in all dimensions.

FIDO contains many new sensor technologies to help identify and classify geologic samples. It has a mutli-spectral imaging system that provides high resolution stereo images. A Color Microscopic Imager allows imaging of fine detail. An Infrared Point Spectrometer measures radiation intensities to help characterize any rocks it finds. A rock corer allows inspection inside of rocks.

Tests have been conducted in the Mojave desert because this area closely simulates the terrain found on Mars. "We want to practice this whole procedure, so that by the time we do it on for real on Mars, we're very good at it," said Dr. Eric Baumgartner, a mission engineer for field tests at JPL.

Robots in Space

Although it is not common to think of space probes as robots, they often have all of the same components. One such exploration craft, the Deep Space One, even had a control system which originated in university robotics labs.

The New Millenium Remote Agent (NMRA) was the control system on the Deep Space One space probe which was launched last year. The NMRA uses a reactive controller based on the RAP (or Reactive Action Packages) language. RAP was developed for robotic control by R. James Firby and has become fairly mature over the years. The NMRA executive links a RAP reactive plan execution system with a novel error reasoning system called Livingstone which is able to identify failures and reconfigure the reaction system to avoid them.

The system works by having a planner retrieve high-level goals from the mission database. A scheduler determines a course of action to fulfill the goals and sends a schedule to the reactive executive. The reactive executive uses behavior-based reactive execution to carry out the steps in the schedule. The error recovery system examines the spacecraft status and modifies the execution of tasks as necessary to compensate for problems. The Deep Space One is an example of a complicated robot in space!

A Primer On Battery Usage

by Arthur Ed LeBouthillier

In developing mobile robots, one is often faced with severe limitations on the options available for power sources. Sure, one can consider internal-combustion engines, steam power, solar power and a host of other options, but each of these has unique problems that limits their use to particular environments or to particularly large robots. Batteries are about the only real option available for mobile robots especially with robots for in-the-house usage. Additionally, rechargeable batteries are the only real economical option for hobbyists. It helps to understand how to determine the proper batteries for your application.

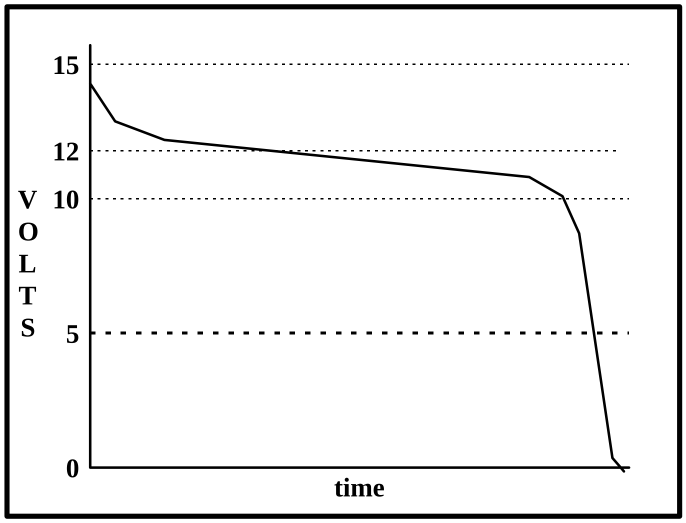

Figure 1 - 12 Volt Battery Discharge

Battery Basics

Batteries use dissimilar materials to create electrical energy from chemical energy. There are two major kinds of batteries: primary and rechargeable. Primary batteries refer to batteries based on non-reversible chemical processes. Rechargeable batteries use reversible chemical processes allowing you to use them time and time again.

There are two major factors to consider when choosing a battery: the voltage and the current rating. The voltage rating is fairly straightforward; a battery is rated as providing a particular voltage for the duration of its charge. In reality, a battery never supplies precisely its rated voltage; rather, it supplies around the rated voltage and the voltage specification is an average over its charge period. At full charge, a battery generally has more than its rated voltage. A 12 Volt lead-acid battery, for example, can supply over 13.8 volts. As it reaches the end of its charge, a battery also supplies less than its rated voltage. Again, for a lead-acid battery it could go below 10.1 volts. However, between its high charge value and its low charge value, most batteries supply a fairly linear drop in voltage over their discharge period. Figure 1 shows a discharge cycle for a typical battery.

So how much time does a battery take to discharge? That is specified by the Amp-Hour rating. What that means is that the battery will put out its specified voltage (on average) for one hour at the specified amps. Therefore a 12 Volt 4 Amp-Hour could theoretically supply 4 amps for one hour; likewise, it could supply one amp for four hours. In reality, you are likely to see less than the maximum rating when you draw higher currents. The specification for a battery is measured over a period of 10 hours and scaled to be equivalent for 1 hour. When measuring a battery designated as 12 Volt 4 Amp-Hour, it was tested by drawing 0.4 Amps for 10 hours. Drawing at higher amperage causes greater losses due to heating and internal resistances. Therefore, to see the rated charge duration for a battery, you shouldn't draw current at rates which are too high. Again, the way to figure out "too high" is to realize that the specification calls for drawing the power over a ten hour period and as you draw current at a rate that would lead to a battery discharging too soon before that, you are approaching "too high" of a current draw and will see less overall power from the battery.

Power and Energy Density

Another real factor in choosing a battery for a mobile robot is how much power or energy it can supply per pound. As a quick refresher, Watt determined that power (in units of Watts) is equal to the voltage (in units volts) times the current (in units

(see Batteries on page 4)

Type	Energy Density W-h/lb	Power Density W/lb	Self-discharge/ month	Cycles to 80%
Pb-acid	14-20	91	5%	200-1000
Ni-Cd	18-23	86	15%	500-1000
NiMH	23-27	82	25%	500-1000
Li+	59	363	5%	1200
Ag-Zn	64-91	45-150	4%	100-250
Ag-Cd	25-43	45-100	4%	300-500
Zn-Air	91-136	36-45		N/A
Al-Air	159	227-272		N/A

Table 1 - Comparison of battery technologies

of amps). The Energy Density refers to the number of watt-hours that a battery can supply per pound. Table 1 shows a list of different battery technologies and their energy densities.

As can be seen, lead-acid batteries (Pb-acid) have an energy density of 14-20 Watt-hours per pound. Nickel-Cadmium (Ni-Cd) can have a slightly higher energy density than lead-acid batteries and Nickel-Metal Hydride (NiMH) batteries can supply even more energy per pound. Conversely, these other battery technologies may not be able to supply as high a current for each pound that lead-acid can.

Listed are also a number of other battery technologies which may be very difficult to use (because of temperature requirements) or expensive. Nonetheless, it is interesting to rate the batteries you might actually plan on using against some alternative “exotic” battery technologies.

One of the more interesting technologies of the exotic battery technologies is the lithium ion (Li+) battery because it provides very high energy and power densities and is reasonably easy to use (although it requires a high-quality charger). A Li+ battery can hold almost twice the energy of lead-acid batteries for the same weight.

Even more exciting than the Li+ battery is the aluminum-air battery (Al-air). It provides extremely high energy density as well as high power density. Yardney, a manufacturer of these batteries, has one model which is only about 13 pounds but which is able to supply 3.0 kW-h at 12 volts. This is equivalent to having a 250 amp-hour battery at 12 volts! Of course, such exotic battery technology costs much more than the equivalent lead-acid battery but for some applications, this cost can be justified. One disadvantage of the Al-Air batteries is that they are not electrically rechargeable, because the anode (the aluminum plates) must be replaced to perform the recharge function. Therefore, they are more like primary cells with replaceable plates.

Summary

A battery is a chemical device able to supply a certain amount of current at a voltage for a duration. One can determine his battery needs by finding out the required operational parameters and then choosing a battery type that will fulfill this need.

<u>Source</u>

Dowling, Kevin, <u>Power Sources for Small Robots</u>, CMU-RI-TR-97-02

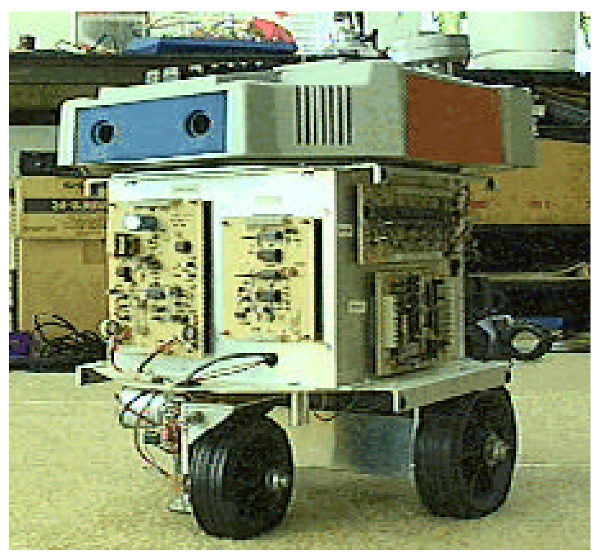

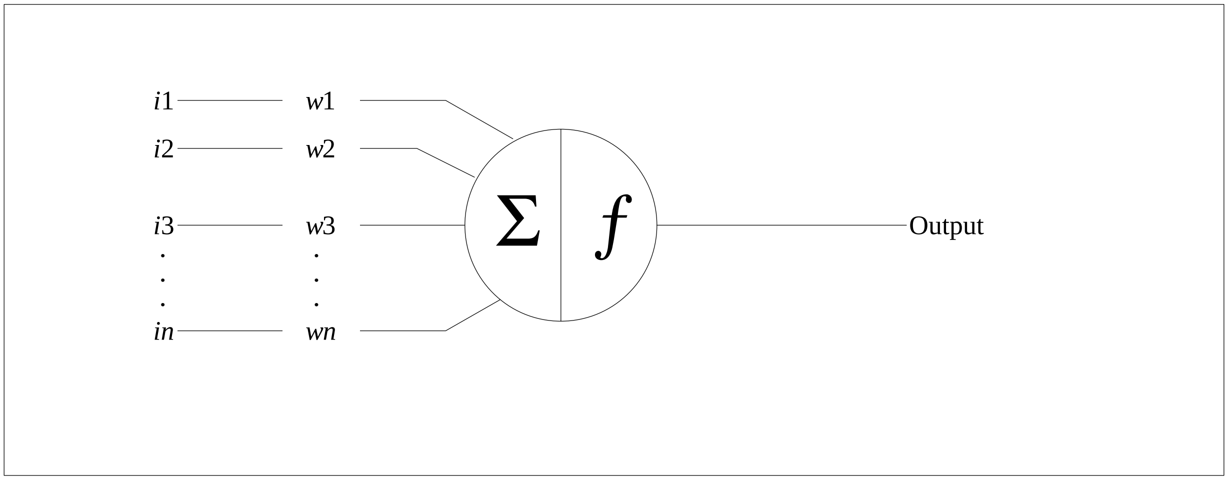

Figure 2 - a simple model of a neuron

of propagation delays, some kinds of processing can occur in the dendrites as coinciding signals add to or subtract from each other.

The Neural Net Model

After studying the way that neurons work, various models have been created to explain them. Some of them are very complicated, such as cable or compartment models, whereas others are simpler. One of the simpler models is the Neural Net model (as opposed to Neuronal Nets). The neural net model simulates each cell as a simple summation of the inputs times a weight going through a non-linear function called a sigmoid function.

Figure 2 illustrates this model. There are the inputs (*i*1,*i*2,…,*in*) with weight values (*w*1,*w*2,…,*wn*), a summation of the inputs (Σ) and a sigmoid function (*f*). Mathematically, they are related by the equation:

output = f(*i*1 * *w*1 + *i*2 * *w*2 + … + *in* * *wn*)

The input values can range from 0 to 1, the weight values range from -1 to 1 and the output varies from 0 to 1. Excitatory inputs use positive valued weights, inhibitory inputs use negative valued weights. Notice that this model does not include any representation of the delay inherent in a real neuron; however, discrete delays can be modeled by using a single-input neuron with a unity weight. This will provide a discrete delay equal to the time it takes to process one neuron.

The Sigmoid function

So far, the model looks pretty simple: all we need to do is multiply each input with its associated weight value, summing all of these individual results together. Once we have this sum, we then need to run it through the sigmoid function. There are a number of different sigmoid functions used. Whichever function is used, they all have the same characteristic. They generally exhibit an "s" shaped function as demonstrated in figure 3.

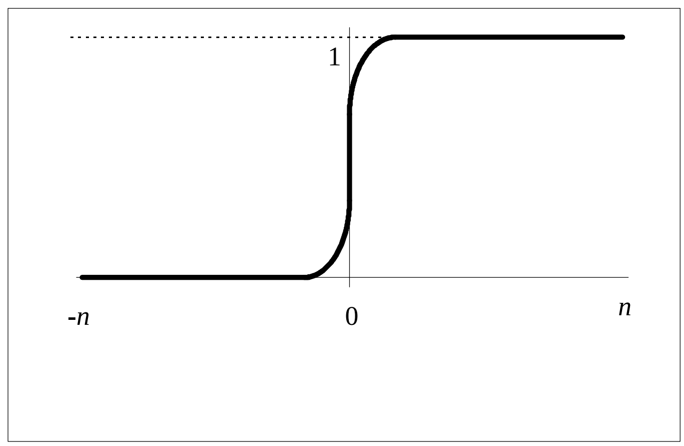

Figure 3 - the shape of a sigmoid function

Figure 3 illustrates a function which stays low throughout a wide portion of its range and which rather suddenly, but continuously, jumps to a positive 1 value. This extreme non-linearity at the transition point models the all-or-none characteristic of real neurons. If the sum of the inputs times their weights reaches the transitition value, then the output quickly jumps up to the full activation value. As was said earlier, a number of sigmoid functions exist. Figure 4 illustrates one typical one.

(see Neural Netl on page 6)

Neural Net from Page 5

$$\text{Output} = \frac{1}{(1 + e^{-sum})}$$

Figure 4 - A typical sigmoid function

This function will provide the sigmoid shape needed for the output. It takes as input the sum of the weights times the inputs and produces an output which varies between 0 and 1. The constant ***e*** is the engineering constant, 2.71828.

The Multi-Layer Perceptron

Now that we have a simple model of a neuron, we can construct more neurons into basic circuits. One of the more important structures that can be created is the Multi-Layer Perceptron (MLP). The multi-layer perceptron has been shown to be adequate for nearly any purpose if you use enough neurons. Figure 5 shows a simple example of a multi-layer perceptron.

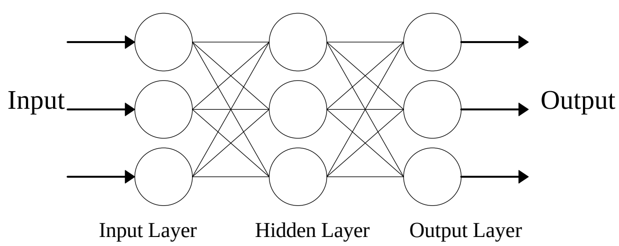

Figure 5 - the multi-layer perceptron

A multi-layer perceptron consists of an input layer, a hidden layer and an output layer. It is able to perform basically any functional relationship between its inputs and outputs. For example, you could use a multi-layer perceptron in optical character recognition; the MLP could be taught to relate the video pixels surrounding a picture of an alphabetic letter to the ascii code representing it.

Implementing the Simple Neural Net Model

Using neural nets once one knows the weight values is simple. The primary operation in the neural net math is a multiply and add; we multiply the weight times the input value and add that to the running sum. In the following code example, we see that the operation is very simple.

```
' ---------------------------------------------
' A simple 3X3 multi-layer perceptron
' This program assumes that the values
' of the weights have already been put
' into the weight arrays

' Get the input values
input_layer(1) = input_1
input_layer(2) = input_2
input_layer(3) = input_3

' calculate the hidden layer  values
for hiddenCell = 1 to 3
  sum = 0
  for weightCnt = 1 to 3
    sum = sum + hidden_weight(hiddenCell,weightCnt) * input_layer(WeightCnt)
  next weightCnt
  hidden(I) = sigmoid(sum)
next hiddenCell

' calculate the output values
for outputCell = 1 to 3
  sum = 0
  for weightCnt = 1 to 3
    sum = sum + output_weight(outputCell,weightCnt) * hidden(weightCnt)
  next weightCnt
  output(outputCell) = sigmoid(sum)
next outputCell

' send the output values
output_1 = output(1)
output_2 = output(2)
output_3 = output(3)
```

This program is quite simple, able to map three input values to three output values. However, as it exists, the program will not work because there are no values in the weights.

Learning Algorithms

How do we get the weight values? This requires a learning algorithm. We must engage in a training session with the neural net whereby we provide the input samples and train the network to get the proper output values. To do this requires an algorithm that compares the actual output with the desired output for a given input and determines the proper values for the weights.

There are numerous learning algorithms, which determine the weights and connections, such as Hebb's learning algorithm, the back-propagation algorithm and others. Each of these algorithms, as part of a learning session, will produce the weight values necessary to make a neural net work properly for its training set and similar inputs. There are also algorithms that are able to learn without being told what their results should be.

Next Month

We'll look at a simple learning algorithm allowing you to teach a neural net to do something useful.

Robotics Society of Southern California

President	Randy Eubanks
Vice President	Henry Arnold
Secretary	Arthur Ed LeBouthillier
Treasurer	Henry Arnold
Past President	Jess Jackson
Member-at-Large	Tom Carrol
Member-at-Large	Pete Cresswell
Member-at-Large	Jerry Burton
Faire Coordinator	Joe McCord
Newsletter Editor	Arthur Ed LeBouthillier

The Robot Builder (TRB) is published monthly by the Robotics Society of Southern California. Membership in the Society is $20.00 per annum and includes a subscription to this newsletter.

Membership applications should be directed to:

Robotics Society of Southern California
Post Office Box 26044
Santa Ana, CA 92799-6044

Manuscripts, drawings and other materials submitted for publication that are to be returned must be accompanied by a stamped, self-addressed envelope or container. However, RSSC is not responsible for unsolicited material.

We accept a wide variety of electronic formats but if you are not sure, submit material in ascii or on paper. Electronic copy should be sent to:

apendragn@earthlink.net

Arthur Ed LeBouthillier - editor

The Robotics Society of Southern California was founded in 1989 as a non-profit experimental robotics group. The goal was to establish a cooperative association among related industries, educational institutions, professionals and particularly robot enthusiasts. Membership in the society is open to all with an interest in this exciting field.

The primary goal of the society is to promote public awareness of the field of experimental robotics and encourage the development of personal and home based robots.

We meet the 2nd Saturday of each month at California State University at Fullerton in the electrical engineering building room EE321, from 12:30 until 3:00.

The RSSC publishes this monthly newsletter, The Robot Builder, that discusses various Society activities, robot construction projects, and other information of interest to its members.

Membership/Renewal Application

Name ____________________

Address ____________________

City ____________________

Home Phone ( ) - Work Phone ( ) -

Annual Membership Dues: ($20) Check #
(includes subscription to The Robot Builder)

Return to: RSSC
POB 26044
Santa Ana CA 92799-6044

How did you hear about RSSC? ____________________

RSSC
POB 26044
Santa Ana CA 92799-6044

Please check your address label to be sure your subscription will not expire!