Article on neural networks

Unedited posts from archives of CSG-L (see INTROCSG.NET):

Date: Thu, 18 May 1995 11:47:14 -0500
Subject: Lobster Neural Network

[From Bruce Abbott (950518.1145 EST)]

The May-June 1995 issue of _American Scientist_ contains several articles of potential interest to PCTers, including the one I'll describe here, entitled "Dynamic Networks of Neurons," authored by Simmers, Meyrand, and Moulins. The topic of the reported research is a "central pattern generator" that produces rhythmic, coordinated ingestion and food-processing movements in the foregut of the lobster. The pyloric network in the somatogastric ganglion consists of 14 neurons, most of which are motor neurons, which are interconnected both via inhibitory synapses _and_ a bidirectional _electrical_ pathway. A sine-wave-like pattern of ion channel openings/closings on the cell membranes of some of the motor neurons produces a rhythmic pattern of polarization/depolarization in these cells that contributes to the production of a regular pattern of impulses to produce a specific sequence of muscle contractions. Normally these outputs generate activities in three separate areas of the foregut for grinding and filtering the food. However, input from "pyloric suppressor neurons" produces a functional change in the way the pyloric cells interact; as a result the three independent activities cease and are replaced by a coordinated action which produces swallowing and motion of the food through the gut. The pyloric system has been shown to operate normally after dissection (when placed in a physiological saline solution without any sensory connections that might supply feedback; thus the system is open-loop. (This makes sense given the "mission" of the system; however, it may be that the switch from grinding/filtering to swallowing/moving is closed-loop, a possibility not touched on in the article.)

What I think is of interest to PCTers in the article is not that the system is an open-loop one, but the way in which its pattern of output can be changed, effectively reorganizing the circuit. The pattern is not only influenced by the suppressor-neuron activity, but also is sensitive to the levels of various neurotransmitters, which can alter such characteristics as the threshold for initiating an action potential and whether there is just a single spike or a sustained depolarization of the neuron the produces a high-frequency succession of impulses. The authors note that understanding of neuron functioning is changing:

> Until rather recently, most neurobiologists considered neurons as logical threshold units that sum incoming signals and then linearly transform the analog input into digital output, or pulses. We now know, however, that neurons are not simply "all-or-nothing" devices that can only be on or off. Instead, they can possess additional bioelectrical properties that have far-reaching consequences for their computational ability and function.

The authors do seem to buy into the notion that coordinated behavior such as walking is produced by "rhythm-generating circuits" (controlling output rather than input), but despite that defect the article is worth reading for the information it provides about how pre-organized systems of neurons can effectively become temporarily reorganized to perform alternate functions, and for its tutorial on current understanding of neuronal function.

Regards, Bruce

Date: Thu, 18 May 1995 22:08:23 -0600
Subject: Re: Lobsters

[From Bill Powers (950518.2020 MDT)]

Bruce Abbott (950518.1145 EST) --

My impression of things like the lobster neural network (there have been others, such as for Aplysia) is that I'm hearing people say "Here's a green dot, and right next to it is a red dot at bearing 120 and a blue dot at bearing 240, and then a row of black dots at approximately equal intervals, and then crossing that row, a row of green dots next to a row of black dots ....". I keep wanting to yell, BUT WHAT IS IT? If you back away just far enough, you see that it's a color half-tone picture of somebody.

Of course it's not that bad in this case, but there's some of the same flavor.

> The pyloric network in the somatogastric ganglion consists of 14 neurons, most of which are motor neurons, which are interconnected both via inhibitory synapses _and_ a bidirectional _electrical_ pathway. A sine-wave-like pattern of ion channel openings/closings on the cell membranes of some of the motor neurons produces a rhythmic pattern of polarization/depolarization in these cells that contributes to the production of a regular pattern of impulses to produce a specific sequence of muscle contractions.

If you think about our models, how much could you convey about how they work if you just described the connections and didn't specify the parameters? Do you think that a listener drawing a diagram from your description would ever tumble to the fact that with a high gain in the loop and the right slowing factor, this system could control something?

For me, the bare recitation of where "impulses" go is almost totally uninformative, especially when the report includes details like calcium channels opening and closing. I want to know the transfer functions -- the relationships of input signals to output signals for each neuron, in terms of frequencies. Maybe the rhythmic signals are being produced by a single-neuron oscillator, but maybe they're produced by a feedback system with parameters set to make it drive itself into oscillations, with neural signals being looked at in terms of frequencies. When you look so closely that all you can see are the individual impulses, it's just like looking at a half-tone picture under a microscope. You can see what is happening in great detail, but you can't see WHY it is happening.

> The pyloric system has been shown to operate normally after dissection (when placed in a physiological saline solution without any sensory connections that might supply feedback; thus the system is open-loop).

Don't forget that every individual function in a control system runs open-loop. If the output function is a variable-frequency variable- amplitude oscillator (something Rick and I are looking into -- now and then -- as a model of controlled repetitive movements), it will produce oscillations when driven by a signal. But the consequences of those oscillations will not be controlled unless there is perceptual feedback reporting those consequences. If you dissected the integrator out of one of our control-system models (or the real system) and stuck signals into its inputs, you would see that the output would be approximately the integral of the input signal. It would be operating "normally". But it sure as hell wouldn't be controlling anything.

> What I think is of interest to PCTers in the article is not that the system is an open-loop one, but the way in which its pattern of output can be changed, effectively reorganizing the circuit. The pattern is not only influenced by the suppressor-neuron activity, but also is sensitive to the levels of various neurotransmitters, which can alter such characteristics as the threshold for initiating an action potential and whether there is just a single spike or a sustained depolarization of the neuron that produces a high-frequency succession of impulses.

The _system_ isn't open-loop; one component of it is.

Control through alteration of system parameters is an interesting subject, which you will see in the model Rick and I are working on if we ever get it going. But it's not reorganization in the technical PCT sense. I am going to be unhappy with reports like the one about altering characteristics of neurons until I see a study that explores the range of behaviors of the neurons over a range of input frequencies of impulses or concentrations of neurotransmitters. I can't believe that the choice is simply between a single spike and a sustained high- frequency output. If that's all you test for, of course, you'll see only those two states. But with the right inputs, you might see that there is a continuous transition between those extremes, and then you might get some idea of what this neuron is doing. If you don't do the experiments in a way that would reveal the continuous transition if it exists, you'll never know if your observations are real observations or just samples at the extremes.

> The authors do seem to buy into the notion that coordinated behavior such as walking is produced by "rhythm-generating circuits" (controlling output rather than input)

I think that a walking-control system is going to need an output function that is an oscillator -- a rhythm-generating circuit. That will produce the basic changes in leg position. But that isn't enough to allow bipedal or other locomotion, because it's open-loop. There's no provision for correcting unpredictable errors. You need a way of varying the amplitude and frequency of the oscillations based on their consequences, such as the effect on foot placement on uneven or tilted ground, and also varied as a way of maintaining balance and changing direction.

Best to all, Bill P.

Date: Fri, 19 May 1995 10:48:12 -0500
Subject: Re: Lobsters

[From Bruce Abbott (950519.1045 EST)]

>Bill Powers (950518.2020 MDT) --

> If you think about our models, how much could you convey about how they work if you just described the connections and didn't specify the parameters? Do you think that a listener drawing a diagram from your description would ever tumble to the fact that with a high gain in the loop and the right slowing factor, this system could control something?

> For me, the bare recitation of where "impulses" go is almost totally uninformative, especially when the report includes details like calcium channels opening and closing. I want to know the transfer functions -- the relationships of input signals to output signals for each neuron, in terms of frequencies. Maybe the rhythmic signals are being produced by a single-neuron oscillator, but maybe they're produced by a feedback system with parameters set to make it drive itself into oscillations, with neural signals being looked at in terms of frequencies. When you look so closely that all you can see are the individual impulses, it's just like looking at a half-tone picture under a microscope. You can see what is happening in great detail, but you can't see WHY it is happening.

My very brief quote from and comments on the Simmers, Meyrand, and Moulins article was only intended to whet the appetite. These researchers have done a very nice job of identifying the input/output relationships of this system, as you will see when you read their article. They have measured population responses in the relevant motor nerves and related these to the activities of the individual neurons within the ganglia.

> Don't forget that every individual function in a control system runs open-loop.

I haven't. Here we are only examining a part of a larger system, the part that generates the outputs which actually produce muscle contractions. It is not a control system, but it may be a part of one.

> The _system_ isn't open-loop; one component of it is.

Not all systems are control systems. By "system" I meant a collection of elements organized to serve some function or functions. For example, the photoreceptors, bipolar cells, amacrine cells, horizontal cells, and ganglion cells of the retina compose a system that converts a pattern of illumination into a pattern of neural impulses. The system described by Simmers et al. is open-loop, although it may be a component of a control system.

> I am going to be unhappy with reports like the one about altering characteristics of neurons until I see a study that explores the range of behaviors of the neurons over a range of input frequencies of impulses or concentrations of neurotransmitters. I can't believe that the choice is simply between a single spike and a sustained high- frequency output. If that's all you test for, of course, you'll see only those two states. But with the right inputs, you might see that there is a continuous transition between those extremes, and then you might get some idea of what this neuron is doing. If you don't do the experiments in a way that would reveal the continuous transition if it exists, you'll never know if your observations are real observations or just samples at the extremes.

Please, read the article! The authors have done these tests. Furthermore, in the system under discussion, it makes perfect sense that it would operate the way it does.

> I think that a walking-control system is going to need an output function that is an oscillator -- a rhythm-generating circuit. That will produce the basic changes in leg position. But that isn't enough to allow bipedal or other locomotion, because it's open-loop. There's no provision for correcting unpredictable errors. You need a way of varying the amplitude and frequency of the oscillations based on their consequences, such as the effect on foot placement on uneven or tilted ground, and also varied as a way of maintaining balance and changing direction.

I'm pleased that you recognize the need for oscillator circuits to orchestrate some patterns of output, and I agree that such circuits do not eliminate the need for control. A particularly well-worked out example of an oscillator circuit in biology is the wing-beating mechanism found in insects (sorry, I don't have the references, but I could find them if you're interested). Flying is, of course, controlled; the oscillator simply supplies the beat frequency.

The mechanism Summers et al. discuss, in which the pattern of output mediated by a particular ganglion changes dramatically depending on the state of its input (active or quiescent), is a biological realization of Ross Ashby's (1956) "machine with input."

Regards, Bruce

Date: Fri, 19 May 1995 18:59:19 -0600
Subject: Re: Lobster article

[From Bill Powers (950519.1700 MDT)]

Going up to Boulder tomorrow, for three days, to see new granddaughter!

Bruce Abbott (950519.1045 EST)

Rick Marken is sending me a copy of the lobster article. Will reserve further comments until I've read it.

Best to all, Bill P.

Date: Wed, 24 May 1995 11:17:44 -0600
Subject: Lobsters

[From Bill Powers (950524.0915 MDT)]

Bruce Abbott (9505xx) --

Thanks to Rick Marken, I finally have the article by Simmer, Meyrand, and Moulins, "Dynamic Networks of Neurons" (the lobster article) and have been through it a few times.

It seems that the main point is that by application of signals from higher centers, functional units of neurons can be temporarily used as if they were larger units. Considering HPCT, this does not come as a great surprise, although it's nice to see the detailed work being done.

It would be nicer if the authors weren't so prone to overinterpretation. Of course that impression may come from the popularization mode of presentation; in the original articles there may be a lot more detail.

One problem I have is with the method of recording data. The traces of Fig. 3 seem to be a combination of impulse recordings with an electrode (or amplifier circuit) response that rectifies and semi-smooths the spikes, so we get a quasi-average-frequency measure superimposed on spike information. There seems to be some high-frequency cross-talk between electrodes which may have nothing to do with function. I think it must be very difficult to get a recording that shows physiologically meaningful variations in single cells with sufficient bandwidth to make sure we're seeing what is happening in the cell instead of in the electronic amplifiers or the surrounding medium. An electrode stuck into the fluids around a cell isn't exactly a high-impedance probe, and who knows what aspect of cell function it is picking up?.

According to the text, the recordings were made with electrodes on the outgoing motor axons, not in the cell-bodies as shown in Fig. 3. These are described as "extracellular" recordings, which means they are subject to influences that may have little to do with the cell functions.

The main problem with this article (if not the original studies) is that only the outputs of the cells are shown, with nothing to indicate the nature of the connections between them. What we seem to have is a neural oscillator, with the two neurons combined into a circuit that involves one cell (PY) firing at rates that vary 90 degrees out of phase with the firing-rate variations in other cells. Mutual inhibition alone will create a flip-flop effect; there must also be time integration to get an oscillator. I don't have the impression that the studies were sufficiently detailed to pick up a time integration effect.

It would be interesting to know what the effect of a bidirectional electrical connection between two cells is. When either one fires, does it make the other fire, too? Obviously you can't have impulses going in both directions literally simultaneously; colliding normal and antidromic impulses should simply cancel, as the mechanism for propagation has been depleted immediately behind each impulse. The only other possibility is that the link acts like a electrical wire, forcing potentials at both ends to be the same.

As to the "bursting" property that is gated on and off by a "modulatory interneuron" in Fig. 6, this looks like a gated one-shot. An initiating electrical spike triggers the one-shot on if the enabling input is present, after which the one-shot maintains itself on until some internal integrator reaches the threshold needed to sustain the firing. Then the one-shot turns off. I would expect that other parameters, like the frequency during the on period and the duration of the on period would also be adjustable by external signals.

Almost any neuron will behave something like this if presented with a large enough input stimulus. If a large jolt of excitatory neurotransmitter is received, the innards of the cell are biased far above the threshold of firing, and there will be "after-firings" for a time depending on how extreme the stimulus was (how many parallel inputs at the same instant) and on the amount of chemical concentration that is used up or diffused away with each firing. Some cells (for example, local negative feedback internuncials leading from a motor neuron's output back to its cell body) have after-firing rates that fall off exponentially to zero after each impulse, leading to negative leaky- integral feedback around the neuron, providing a phase-advance circuit (in terms of frequency of firing).

Everything depends on the parameters. The connections alone tell us only one part of the story.

I was struck by the complete absence of any discussion of the role of sensory signals. The reason may be in the authors' stated theory:

> Animals choose constantly among a wide range of behavioral capabilities: walking, running, fighting, courting, and so on. An animal's nervous system generates each of these activities by turning on a specific network of neurons, which produces a characteristic sequence of electrical signals, called motor output, that instructs muscles to perform particular movements.

Obviously the authors believe that behavior is simply output, with all actions, however complex and however related to unpredictable environmental events, merely being emitted and having fixed effects. So naturally, they concentrate on output functions. If they don't realize that walking, running, fighting, and courting require constant readjustment of outputs in order to maintain the patterns we observe, they will obviously not realize that generating output doesn't explain behavior. I should think that they would at least mention the fact that sensory feedback from food in the esophagus, stomach, and pylorus must have something to do with the way these functions are used.

As to extending these results to larger units of behavior, there are some severe problems. If simple oscillators were sufficient to explain things like walking, maybe the parallel would hold, but they aren't. To prove this to yourself, just start walking and then slow down and reverse. You can actually freeze the walking at almost any point, reverse it, continue in the same direction, or go back and forth between forward and reverse. Maybe this can be accomplished with an oscillator circuit that has parameter-setting inputs, but it's pretty hard to imagine. The simple image of an oscillator that produces regular "pendular" motions of the legs is obviously insufficient. (I've always thought that the image of "pendular" movements is singularly inappropriate, as the pendulums are actually upside-down with the fixed point in contact with the ground).

Nice observations on fly behavior

Best to all, Bill P.

Date: Fri, 26 May 1995 11:34:33 -0500
Subject: Lobster Tales; A Leg to Stand On

[From Bruce Abbott (950526.1100 EST)]

>Bill Powers (950524.0915 MDT)

> Thanks to Rick Marken, I finally have the article by Simmer, Meyrand, and Moulins, "Dynamic Networks of Neurons" (the lobster article) and have been through it a few times.

I got my copy off the rack at Reader's World. For those who might like a copy, it's the current issue of _American Scientist_ (May-June) and should still be available.

> According to the text, the recordings were made with electrodes on the outgoing motor axons, not in the cell-bodies as shown in Fig. 3. These are described as "extracellular" recordings, which means they are subject to influences that may have little to do with the cell functions.

It's only a minor point, but I think you've misunderstood the text. The top trace of Figure 3 provides the extracellular recording of motor-nerve bundle activity--a population response. The bottom two traces are from intracellular probes, as shown in the figure, and give the outputs of two individual cells, one PD and one LP.

> I was struck by the complete absence of any discussion of the role of sensory signals. The reason may be in the authors' stated theory: . . .

> Obviously the authors believe that behavior is simply output, with all actions, however complex and however related to unpredictable environmental events, merely being emitted and having fixed effects. So naturally, they concentrate on output functions.

That was my impression, too. Another reason may relate to the article's focus on the "dynamic rewiring" of the somatogastric network to change its mode of action from grinding/filtering to swallowing/transporting. A discussion of sensory inputs and the larger system within which the somatogastric network functions may have been viewed as outside the scope of the article.

> As to extending these results to larger units of behavior, there are some severe problems. If simple oscillators were sufficient to explain things like walking, maybe the parallel would hold, but they aren't. To prove this to yourself, just start walking and then slow down and reverse. You can actually freeze the walking at almost any point, reverse it, continue in the same direction, or go back and forth between forward and reverse. Maybe this can be accomplished with an oscillator circuit that has parameter-setting inputs, but it's pretty hard to imagine. The simple image of an oscillator that produces regular "pendular" motions of the legs is obviously insufficient.

I was more interested in the demonstration that different output functions could be carried out by the same neural system as a function of input state than in the neural oscillator function itself. Very economical. However, neural oscillators are no doubt found here and there in the human nervous system, probably most often to produce rhythmic, patterned muscle contraction for vegetative systems such as digestion, breathing, and heartbeat. Such a circuit might also be involved in producing circadian rhythms, which include the sleep/wake cycle.

> (I've always thought that the image of "pendular" movements is singularly inappropriate, as the pendulums are actually upside-down with the fixed point in contact with the ground).

The pendulum is the free leg, not the one you're standing on. It has a natural period of oscillation, so that the most comfortable speed of locomotion, requiring the least effort, is at the resonant frequency of forced oscillation.

> Nice observations on fly behavior.

Thanks, glad you liked them.

Regards, Bruce

Date: Fri, 26 May 1995 12:12:44 -0600
Subject: Lobsters

[From Bill Powers (950526.1100 MDT)]

[Bruce Abbott (950526.1100 EST)]

> I was more interested in the demonstration that different output functions could be carried out by the same neural system as a function of input state than in the neural oscillator function itself.

From what I got out of the article, it seemed that the individual oscillators just kept working the same way; what changed was their synchronization. or whether some were turned on or off. This is more like higher-level control achieved through actions on lower systems of fixed organization.

The idea of using a set of lower-order components to achieve different functions when driven by higher-order systems is surely nothing new to us, is it? I use the same hand to scratch that I use to write. What's the big deal?

What I meant by saying this sort of neurological analysis is a crock was not that a competent analysis is a crock, but that an analysis that simply invokes a new causal signal every time a new effect is needed is a crock. To do a competent analysis of a neural system, you have to know the input-output laws that govern each component, and derive the behavior of the whole system from the interactions of the components. Same problem as deducing what an electronic circuit does from reading the circuit diagram. Unless you know the laws governing resistors, capacitors, inductors, transistors, transducers, transformers, and so forth, you are extremely unlikely to guess what the circuit really does, even if you can identify input and output signals.

And as Rick pointed out, even if we did understand the neural circuit in this sense, we still would have to know the rest of the loop to make any predictions about what such a system would actually do.

In B:CP I laid out a few basic components and discussed what their properties would be. Such components could be hooked together in many different ways (as with electronic components) to achieve complex functions. Of course it would be even better to have a set of components which have been derived from the actual properties of neurons.

Best, Bill P.