

Re: Finding useful functions– part 1

Source: <http://sci.tech–archive.net/Archive/sci.cognitive/2004–10/0801.html>

From: David Longley (*David_at_longley.demon.co.uk*)

Date: 10/28/04

Date: Thu, 28 Oct 2004 15:54:30 +0100

In article <nk6gd.15541\$rs5.865264@news20.bellglobal.com>, Wolf Kirchmeir <wwolfkir@sympatico.ca> writes
>Bill Modlin wrote:
>
>[...]
>> *Let's start over.*
>
>OK.
>
>> *You describe pretty much the same situation I see, in which a cell
>> "learns"... i.e. changes its various operating parameters, thus
>> changing its functional mapping between inputs and outputs, based on
>> purely local signals. These local signals include its own firings,
>> activations of its synapses by other cells firing, and changes in
>> various chemical concentrations in which it is immersed. As you
>> say, none of these inputs are labelled.*
>
>Cells don't learn. That's a hierarchy error – its the network that
>learns, maybe. Certainly the network of networks that we call, say, the
>visual cortex learns.
>
>The cell's responses to inputs are modified by some inputs, yes, but
>that is not the same as learning. It's analogous to the content of a
>RAM location being modified by some inputs. I would;nt say the memory
>location learned anything – its electrical charge changed, is all.
>
>> *It seems obvious to me that under these circumstances, if there is
>> to be any systematic rule or principle guiding the way the cell
>> changes its response function, it must be formulated in terms of the
>> signals accessible to the cell, with no reference to any possible
>> remote and indirect consequences of those changes.*
>
>Yes.
>
>> *The "adjustment
>> rules" can depend on the strength and frequency of these local
>> signals and on the timing relationships among them, and that pretty
>> much exhausts the available possibilities.*

>
>*In a natural system, it also includes the chemistry of the surrounding
>medium, which will modify the way the signals act on and in the cell.
>That's a crucial fact, IMO. I.e., "chemical messengers" will promote or
>inhibit the transmission of signals across the synaptic gaps. Since
>these messenger molecules are emitted by other cells, including
>non–neural ones, the picture is much more complex. I haven't a
>conceptual handle on this, certainly not in terms of message content,
>signal labelling, etc etc etc.*
>
>> *This also seems to be the position you are taking. Which confuses
>> me, since on other occasions you seem to argue for a stance much
>> like Glen's, where all "learning" is caused by remote behavioral
>> contingencies.*
>
>*I've already said that cells don't learn. In any case, figuring out how
>neural nets' behaviours/functions change doesn't refute the position
>that such changes are initiated by "remote external contingencies".*
>
>> *Let me try to pose an unambiguous example of the conflict.
>> A pigeon can be trained to discriminate pictures containing trucks
>> from other pictures lacking trucks.
>> This is done by selectively reinforcing some behavior (pecking a
>> button?) in the presence of the truck pictures, and not in the
>> presence of others.
>> At this level of description, this is a supervised learning process,
>> driven by an experimenter–enforced correlation between rewards and
>> behaviors. The rewards are contingent on the production of the
>> right behavior under the right conditions, and the pigeon contains
>> mechanisms to adjust its behavior to maximize rewards.*
>
>*I see no reason to talk about "supervised" learning processes, since
>that word smuggles in the experimenter's intentions. The pigeon will
>learn in exactly the same in nature, the only difference being that
>random behaviours will be reinforced rather than pre–selected ones. So
>what?*
>
>*The mechanisms that "adjust the pigeon's behaviour" include the
>cellular changes that you seem to think exemplify some other kind of
>learning.*
>
>> *That's fine, so far as it goes.
>> But when I look at discriminating that class of pictures so that it
>> can be recognized as a condition for the rewarded behavior, I see a
>> pretty complicated process. There are billions of cells computing
>> functions of whatever inputs they have access to, responding to all
>> sorts of "features" at dozens of levels, bringing together
>> information from many areas of the picture, to eventually reach a
>> level at which there is a signal of some sort that indicates whether
>> or not there is a truck somewhere in the picture.*
>

>
>So the process is complicated. So what? When I watch a rainstorm, I
>see billions of raindrops, millions of turbulence cells, etc. The
>process seems pretty complicated. The net result is still that things get
>very wet.
>
>> That truck–signal is correlated with the rewards and the behavior,
>> so it makes sense at least at a handwaving level that a supervised
>> learning process could incorporate it, and produce the behavioral
>> modifications that we observe.
>> But most of those intermediate signals in the long path from retina
>> to truck–signal are not correlated with anything in the high level
>> description of the experiment. They aren't correlated with trucks,
>> or rewards, or pecking, and therefore could not have been shaped by
>> any of those things.
>
>Yes, that's true, but why should they be?
>
>An analogous problem: how do the hundreds or thousands of fish in a
>school of fish all "know how to change direction? They don't. Each fish
>knows that the immediately surrounding fish are coming closer or
>getting further away, so it adjusts its direction and speed to maintain
>certain distance. That's a fine example of "local correlations", and
>IMO is the way one must think about it. It doesn't matter where the
>change in direction originates (a few fish see a shark, and change
>direction) – the fish inside the school don't get the message "Shark
>nearby, get out of the way". They get only messages about changing
>distance between themselves, and that's what they respond to.
>
>Let the fish be signals moving through a NN, let the changing distances
>be variations in signals passing between neurons, and let the response
>of the fish be the cellular/synaptic changes. Then the school's change
>of direction is the result of different NN functions. The analogy is
>good enough to clarify the concept, IMO.
>
>> To me it seems obvious that they must be shaped by local rules
>> involving relationships among signals accessible to the cell.
>> Specifically they cannot depend on producing some effect in the
>> external environment and reacting to contingent results of that
>> effect. Their connections to the external environment are extremely
>> indirect, there is no way that any correlation mediated by external
>> contingencies could be communicated to them.
>
>It doesn't matter how indirect the connections to the external
>environment are. See the school–of–fish analogy.
>
>> But Glen seems to say that discriminations must be learned as a
>> result of behavioral contingencies. For example, his response to my
>> original post was:
>
>No, Glen says discriminations of environmental contingencies cause

>learning. You got it wrong, hence our confusion about the relationship
>between environment and learning. The shark sets off a change of
>direction in a few fish, and that change propagates through the whole
>school. Not the perception of the shark, please note. What we see is
>"the school of fish changed direction and escaped the shark." But the
>10024th fish doesn't know there is a shark nearby.
>
>Just so, perception of a red light and the pecking of a key etc set off
>changes in the pigeon that we see as "The pigeon has learned to peck a
>key to get food whenever it sees a red light." But the pigeon's retinal
>cells don't know anything about the food, the pigeon's motor cortex
>know nothing about a red light, and so on. Of course a pigeon is a
>more complex system than a school of fish, but the principle of local
>actions resulting in global behaviours applies equally well.
>
>>>What is important in sensation and perception is that
>>>movement of an animal (or, more specifically, of its
>>>receptors) has consequences. When we sweep our eyes
>>>over a patch of red, there are changes in stimulation –
>>>such movement/consequence contingencies are at the heart
>>>of learning to perceive the world.
>> But this would seem to imply that each of those billions of cells
>> involved in discriminating trucks somehow "knows" that we moved our
>> eyes, and can correlate this with the changes in the other signals
>> it has access to, a claim which I find incredible.
>
>It implies no such thing. But your faulty reasoning does lead to an
>incredible conclusion. IMO you should examine your assumptions,
>specifically the one that cells "know" and "learn."
>
>> Certainly the signals a cell can see generally originate in
>> environmental stimuli. Cells learn their functions from
>> relationships observable in those signals, so all the learning is in
>> a sense ultimately traceable to the environment.
>
>Cells don't learn!!!!!! Get rid of that idea, and things will be much
>clearer. Cells don't learn functions – they have functions. Some of
>those functions depend on the uptake of extra-cellular chemicals, for
>example. But performing those functions when those chemicals are
>present isn't learning, it's just what the cell does.
>
>> But the specific
>> relationships observable at any point in the network are heavily
>> dependent on other functional transforms which generated those
>> signals, and are seldom directly mappable to any particular
>> behavioral contingency identifiable at external levels of
>> description.
>
>I think you are confused about what you mean by a "specific
>relationship observable at any point in the network". I certainly am
>confused about what you mean. Do you mean the likelihood that a signal

>will cross that point? Do you mean the number of other points it's
>connected to? Do you mean the kind of signal present at that point?
>(see above about messenger molecules originating from other parts of
>the organism.) Do you mean the strength of the signal, or its
>frequency if repeated, etc? Etc. Or do you mean the topology of the
>network, both local and global?
>
>> Where do you stand on this matter?
>
>See above. IMO, you have serious gaps in factual knowledge (certainly
>more than I do, and I have a lot), and your conceptualisation of the
>hierarchy of networks and hence of processes within them is vague,
>ambiguous, and error ridden. Eg, you want to locate learning in cells,
>which is like locating a car's motion in the valve train.
>

He's been told this many times before, by both myself and Glen. For some reason, like all too many others in our species he seems to think he knows better. He hasn't told us why he thinks he knows better, and he certainly hasn't shown that he understand the problems that those who have worked on behavioural plasticity ("learning") over the past 70 years or so within the EAB have been addressing. Instead he looks to the literature where he finds comfort from other misguided nitwits who don't actually make any practical headway either. The reality is that a rational person would look to the EAB first to see where their half-baked, (and therefore probably errant) "ideas" have been tacitly imbibed from (through the usual dynamics of our public verbal behaviour). That this advice or point is still so widely argued against (or just ignored) by so many here and elsewhere is an interesting observation about the resistance to behavioural plasticity itself, which is of course just grist to my mill and why I bother to post anything here at all.

Good look at bringing Bill to his senses. I only hope you get better remuneration than Glen and I have – as to date, he's just come back and peddled more recalcitrance in slightly more veiled (albeit still obnoxious) way than some of the other more coarse reprobates here.

--

David Longley