

Re: Finding useful functions– part 1

Source: <http://sci.tech–archive.net/Archive/sci.cognitive/2004–10/0678.html>

From: Bill Modlin (*modlin1_at_metrocast.net*)

Date: 10/26/04

Date: Mon, 25 Oct 2004 22:49:11 -0700

"patty" <pattyNO@SPAMicyberspace.net> wrote in message
news:E59fd.303716\$MQ5.274103@attbi_s52...

> *Bill Modlin wrote:*

> > *Our brains have innate structure tailored by evolutionary
processes*

> > *over a long period of time. This structure performs functions
that*

> > *contribute to our behavior in ways that somewhere along the line*

> > *probably helped individuals to survive, or at least didn't hurt.*

> >

> > *Many of those functions are not fully determined by genetics
alone.*

> > *There is an innate framework, but details are filled in by
processes*

> > *of conditioning and association, and to some degree the
framework*

> > *itself is mutable if environmental conditions differ
sufficiently*

> > *from those for which it evolved. There are few sharp lines
between*

> > *innate and acquired neural function.*

> >

> > *Feature discrimination in the early visual system is sometimes*

> > *called innate. Certainly it is innate that the cells grow into*

> > *layers of tissue appropriate for performing useful feature*

> > *discriminations. However, it seems the specific connections and*

> > *weights to implement particular discriminations get filled in by*

> > *adaptation to correlations in the ensemble of signals flowing
from*

> > *the retina. For example, we can change the distribution of*

> > *particular detectors dramatically by raising a cat in an*

> > *abnormal*

> > *visual environment. It seems cells are not so much genetically*

> > *determined to perform specific discriminations, as that they*

> > *acquire*

> > *discrimination functions appropriate to the signals they
encounter*

> > *in their genetically determined position in the network.*
> >
> > *There are places where neural projections bring together signals*
> > *originating from corresponding points in the left and right*
eyes.
> > *This allows merging both images to fill in details missing from*
one
> > *or the other, estimating depth from discrepancies in the two*
images,
> > *and so on. There is genetic direction to cause axonal*
projections
> > *carrying signals from one eye to grow toward the normally*
expected
> > *locations of the corresponding signal paths from the other. But*
> > *(from experiments on Xenopus frogs) if one eye is surgically*
rotated
> > *before the connections are formed, so that the locations of*
> > *correlated signals are altered, we see the projections grow*
first
> > *toward the normal target location, then veer off sharply to*
connect
> > *with the very different cells now in position to be correlated.*
> >
> > *Many topographic maps can be found in the brain, so that for*
example
> > *neighboring sections of neural tissue are excited by stimuli*
from
> > *adjacent sections of skin. One might imagine a fixed wiring*
scheme
> > *under genetic control to hook up these maps, but when we*
surgically
> > *swap small patches of skin the connections change to preserve*
the
> > *mapping. It takes some time, but after a while we find that the*
> > *moved sensors now activate sections of the remote neural map*
that
> > *correspond to their new positions.*
> >
> > *A reasonable interpretation is that the "wiring" of neural*
circuitry
> > *is only loosely determined by a genetic blueprint. Most of the*
> > *actual connections (and therefore the functions performed) are*
> > *established as a result of correlations between the activities*
of
> > *potentially connected cells. Not only are the initial*
connections
> > *determined by correlations, but even after a stable connection*
> > *pattern is established, the connections will change if the*
> > *correlations change.*
> >
> > *From the viewpoint of a single cell, it strengthens connections*

to

> > *others correlated with its own activity and weakens others, much as*

> > *postulated by Hebb so many years ago. While direct observation of*

> > *such changes in individual active synapses is still difficult, we*

> > *can observe at least one related mechanism in widespread use.*

Cells

> > *in a child's brain sprout huge dendritic trees and eventually make*

> > *something like 200,000 synaptic connections. By adulthood these are*

> > *trimmed back to an average of 10 to 20 thousand. The only plausible*

> > *explanation for this of which I am aware is that the surviving*

> > *connections are those that showed correlation with the activity of*

> > *the cell. Uncorrelated connections simply drop out of the picture.*

> >

> > *Overall, the point is that the functions computed by cells in the*

> > *brain are largely determined by the correlations encountered in the*

> > *signals accessible to the cell, rather than by genetic control.*

> >

> > *This is learning or conditioning, but it is not the kind of*

> > *feedback–driven learning that is usually intended when one speaks of*

> > *operant conditioning. This sort of learning does not depend on*

> > *consequences of the output of the function, and would occur even if*

> > *the output were not connected to anything else and could therefore*

> > *have no consequences extending beyond the cell doing the learning.*

> >

>

> *What evidence do you have that this happens at all ?*

>

> *patty*

I don't understand the question, Patty. The post is a listing of well-known experimental observations of functional changes occurring under circumstances in which the only information available to drive the change is that provided by correlations accessible locally, circumstances which preclude the possibility that behavioral contingency is involved. That seems to me reasonable evidence that it does occur... we can see it happening.

> *From another angle, one might take it as relevant evidence that we can construct a simple neural model, adjust the weights by some sort of Hebbian rule, and find that useful functions are indeed computed. So not only do we see it happening in vivo, we have independent verification that the interpretation of our observations is plausible... mechanisms based on the principle we suppose to be relevant do actually perform as expected, even when stripped down to the barest essentials.*

What sort of evidence are you asking for?

Bill

>
>> *From an evolutionary perspective, such learning mechanisms exist because they do indeed often have useful behavioral consequences.*
>> *But the evolutionary connection is between the learning mechanisms and ensembles of behavior, not between the individual functions learned and specific contingencies associated with those functions.*
>>
>> -----
>>
>> *None of the above should be taken as suggesting that other sorts of learning can be ignored. To implement AI we will require an understanding of many facets of adaptive behavior, including the operant conditioning or reinforcement learning that has been the sole focus of certain vocal participants in CAP.*
>>
>> *But I do suggest that these correlation–driven "unsupervised" mechanisms provide a critically important underpinning for other learning paradigms, that they are necessary parts of an explanation of how all our behavior–generating mechanisms actually work.*
>>
>> *<to be continued in further posts>*
>>
>> *Bill Modlin*
>>
>>
>>
>>
>>