

Re: Finding useful functions– part 1

Source: <http://sci.tech–archive.net/Archive/sci.cognitive/2004–10/0887.html>

From: Wolf Kirchmeir (wwolfkir_at_sympatico.ca)

Date: 10/30/04

Date: Sat, 30 Oct 2004 12:28:10 -0400

Bill Modlin wrote:

>> "Wolf Kirchmeir" <wwolfkir@sympatico.ca> wrote in message
>> news:1GSfd.11528\$Qs6.1183648@news20.bellglobal.com...
>
>
>>> Why should any external feedback be "supervised"? Is there a
>
> difference
>
>>> between supervised and unsupervised external feedback?
>
>
> Just noticed this question that I must have missed when originally
> posted...
>
> If there is external feedback, that constitutes supervision.
> "Supervised" means "relies on external feedback". As opposed to
> "unsupervised" where there is no feedback available.

This sounds like self–organisation to me.

> An unsupervised process still has inputs from external sources.
> It's just that none of those inputs are "feedback", none of them are
> conditional on or functionally derived from its prior outputs.

In other words, if a cat wanders about the yard and sees a bird it hasn't flushed out hiding after hearing it, and responds, then that's an example of an unsupervised process? Whereas if it hears the bird, and flushes it out of hiding, then seeing the bird is a supervised process?

> An unsupervised process must make its own decisions about what sort
> of outputs to produce, usually based on trying to establish some
> sort of desirable statistical relationship to its inputs. This may
> involve local "feedback loops" of various kinds, in which for
> example the last output is used as input to some sort of parameter
> adjustment function. But the point is that there is no external
> process which evaluates the output and feeds back information about

- > *its validity or appropriateness. There is no way for an*
- > *unsupervised process to evaluate the external effect of its outputs.*
- > *It is operating "open loop", blind to the external consequences of*
- > *its actions.*

Since AFAIK the process you describe is exactly what happens when a NN is adjusted as a consequence of some feedback, I fail to see any real difference. IOW, the NN can't tell whether an external input is fed back or not. It's just external input. Besides, you are describing a process that depends on local (neuron-to-neuron) correlations, which cannot, by your own admission even, "know" anything about any external source if inputs. IOW, this process occurs whether or not it's initiated and subsequently "adjusted" by external inputs. So it seems to me that "supervised" and "unsupervised" refer to different levels of the same process. It looks very much that "supervised" means "global", and "unsupervised" means "local." If that's not the intention, the concept is confused, to put it gently.

- > *There are a lot of things that need feedback to succeed. You can't*
- > *use an unsupervised process to learn how to do anything in*
- > *particular, for that you need some feedback.*

Not true, see below: repetition is sometimes enough.

- > *Feedback is such a*
- > *powerful tool, and goal directed processes are so central to our*
- > *thinking about "intelligence" and intelligent behavior, that it*
- > *might seem there is no place in the picture for unsupervised*
- > *processes.*

IMO this is a red herring. Lets leave intelligence out of it for the moment.

- > *But unsupervised processes can be very useful in organising inputs*
- > *and sorting them into categories, activities that depend only on the*
- > *data itself, not on the eventual use of the data for any particular*
- > *purpose.*

What data? Where does the NN get them from? It's at this point that the distinction seems to me to get really confusing. If the NN is operating on the data (eg, sorting it), then the data are external to the NN. The NN accepts them as input and sends the categories (or sorted data or whatever) as output. Whether the data are derived from external feedback seems to me to matter not at all – data are data.

Or are you positing some "internal data"? That's a concept borrowed from computers, and even in that realm it's confusing. All data are external to the algorithm that processes them, whether they are written into the program code or copied from some other source makes no difference whatever.

I can see using the algorithm code itself as data, i.e., changing it. Is that what you mean? But even in that case, the algorithm must treat the

data as external – it must either process a copy of itself, or it must be partitioned so that subroutines within it can treat other subroutines as data. Either way, the data processed is external to the algorithm.

Or do you mean the properties of the elements and connections of the NN? If so, we are back where we started: local correlations between elements of the NN, at which level I can see no difference between supervised and unsupervised as described by you and others.

So my critique amounts to this: un/supervision is a confused concept.

> *Imagine that someone gives you a huge box of miscellaneous strange
> objects, all jumbled together. Even if you don't know what any of
> them are or what you might do with them, a reasonable first thing to
> do might be to sort them into piles of similar-looking things, to
> see how many different kinds there are, and how many of each. And
> then you might start looking for bits that seem to fit together
> nicely, and put those next to each other, or assemble them and look
> for possible larger assemblies.... this is unsupervised activity.*
>
> *Then a supervisor comes along and explains what is to be
> accomplished and starts teaching you how to use them. Or just
> starts giving encouraging smiles when he sees you doing something
> good, so that you eventually figure it out for yourself.*
>
> *But no matter what it turns out you can do with them eventually, you
> are almost certainly better prepared for it by having gone through
> that unsupervised stage of organizing the material first and
> learning what you have to work with, rather than trying to dig
> things out of an unfamiliar jumble when suddenly confronted with a
> supervised task.*

This example is one of global functioning. Moreover, I must already have been trained in sorting in order to do at least this much. When and how was that training done? Or was I just made that way?

Also, "unsupervised" seems to include the notion of self-organisation, as mentioned earlier.

And if this is how it is supposed to work *_within_* a NN, then the neurons must themselves be a kind of network. That's an idea that I've been mulling over the last couple of days, but if it turns out to make sense, it will probably make nonsense of un/supervised.

The discussion of NNs in terms of un/supervised learning, statistical patterns in the data, local correlations, etc, ignores what to my mind is at least equally important: the topology of the connections themselves.

The discussion also ignores what EAB has in fact demonstrated, that external feedback is *_not_* required to produce learning. Mere repetition of external stimuli is enough. EG, make a NN that will produce a

specified output, eg, any given number of repeated outputs from a single input. (That's purely a matter of architecture.)

Now add features such as a second input source, activation potential changes at the connections, and so on. We now have a network that can learn to produce the output from a different range of inputs than I started with. That is, suppose the second input source is connected through inactive nodes. Combine inputs from sources through those nodes until the relevant connections are activated as desired, and then the network will respond the same way to the second input as to the first. Note that this training does not require feedback, just repeated inputs from two sources. But I can't do this unless the architecture of the NN allows it.

Now, if I understand you right, this is an example of unsupervised learning. Yet because it depends on external inputs it's supervised. Or is it both at the same time? Why?

Now add a small complication: two NNs A and B. Let A contain some unactivated connections. Let A initially respond only to external inputs, as above. Let B be an oscillator that produces outputs at intervals. Let B make a connection with A through some inactive node. Let A respond to the external input. Let some of those inputs (depending on relative timing) combine with inputs from B through the inactive node in A. Increase the sensitivity of that node with each repetition of the combined inputs, as above. At some point, A will begin to produce outputs in time with the cyclic outputs of B. — Here too, there has been no feedback, yet external inputs are crucial to the changes in A's behaviour, since without them the inputs from B could have no effect.

Is this supervised learning or not? Why?