Re: Neural netss (was Re: death of the mind.)

Source: http://sci.tech−archive.net/Archive/sci.cognitive/2004−08/0626.html

---

From: dan michaels (*feedbackdroids_at_yahoo.com*)
Date: 08/27/04

```
Date: 27 Aug 2004 12:37:39 −0700
```

Wolf Kirchmeir <wwolfkir@sympatico.ca> wrote in message
news:<_VGXc.30806$DG.1600198@news20.bellglobal.com>...

> The phrase may in fact mislead, since it implies a comparison (and the
> comparison has been explicitly invoked by some contributors to this
> thread.) But I don't see any comparison happening − there is AFAIK no
> "comparator module" that takes input from the "perception module" and
> the "face recognition module" and outputs "Yup, that's Johnny."

"... AFAIK ..." ... as far as you know. How can you possibly say there
is no compararator module, as far as you know? You're just making an
assumption.
===============

The fact
> is that information about edges, etc is produced in the retinal layer,
> that this information passes to various NNs whose output goes to other
> NNs, etc until the final output is "That's Johnny." At each layer, the
> information is reorganised − that is, the inputs to the next layer have
> a different pattern. Note that the face recognition module is not as
> simple as the name impies − some people who have had damage in that area
> can still tell they are looking at a face, they just can't tell who it
> is. Others can't even tell they are looking at a face. Since the damage
> is to the NN, the "reference state" can only be the organisation of the
> NN, as I said above. Again, now what? The organisation of the NN
> determines how it operates, is all. Calling it a reference state doesn't
> help us understand how it's organised.
>

In fact, given the anatomical organization of the visual cortex, with
its 30+ visual areas and 1100+ feedforward and feedback pathways,
there are roughly 1100/2 = 550 pathways which can feed back
"comparison" information. That's what basic FB is − comparison between
previous output and current inputs/state/etc activities.

This whole thing makes sense, regards the cortex, if you think about
the 30+ areas as being mainly filters and comparison areas, with their
inputs coming both from the retina/thalamus/VC directly and also as

feedback signals from the other 29 areas, plus signals coming back
from the "higher−level association" areas in the cortex.

All these signals can be either excitatory, in which case the signals can produce something akin to a resonance/reinforcement effect, or they can be inhibitory, in which case they produce something akin to the comparison that occurs in negative−FB systems, where errors are reduced by the recirculating FB. IOW, a signal coming in directly or via FB can either reinforce or diminish the on−going level of activity. If that activity happens to be a direct "image" input from the retina [ie, spacially organized via retinotopic mapping], then the FB signals can either reinforce or kill the "image" from going any further.

It's undoubtedly true that both forms of activity, positive and negative FB, occur in the cortex at various places, regards those 1100+ pathways mentioned. Individual fiber bundles can easily have both + and − effects, depending upon the layer they terminate at in the cortex. With something this complex, there are all kinds of possibilities. There are billions of possibilities for how the various areas can interact, given all of those interconnect pathways. IOW, from a mechanistic viewpoint, you can have just about anything occurring [although I suspect there are actually a limited number actually in use].

My concept [which is likely to be wrong] is that facial information, per se, is not stored in the facial recognition area of the cortex, rather the FRA is used mainly for filtering of incoming retinal info AND comparison with higher−level info [by virtue of the FB pathways, as described in the previous paragraphs], and that the storage takes place in the synpatic weights in the association areas. IOW, FRA doesn't store anything, rather it's basically organized as a filter for feedforward signals from the retina/thalamus [as are most other of the 30+ visual areas]. Then, the visual areas interact with each other, via those 1100+ pathways, to refine/mutually−reinforce/not−reinforce the activity, and then other inputs come from the association areas so that the signals can be compared to what has been previously stored. If everything matches, that's when you get "Yeah, it's Johnny". That's how the whole thing works. You only get Johnny because billions of neurons and 1000s of pathways factor into the mix.
================

> *Sidebar: The closest computer analog I can see to what happens is bit*
> *masking. However, in a computer, the bit mask is imported from storage −*
> *there is no similar event in the brain that I can see. The NN itself*
> *masks or filters the input; it doesn't import a mask.*
>
> *The reason I keep talking about brain physiology as I understand it is*
> *that I've never been happy with abstract models that ignore the messy*
> *details. "Reference state" is such a concept, since it applies to*

> devices that function differently than NNs. Sure, it's very satisfying
> to recognise some abstract identity (eg, abstractly considered, the
> state of a network can be thought of as stored information), but the
> messy details keep interfering with insight (eg, when the network
> changes state, some or all of the previous information is lost, unlike
> information stored by other means.) Abstract models are useful − they
> obviously help in constructing devices that work−alike some observed
> brain function/system behaviour/etc.
>
> > If you can answer this question, then I have a harder followup. What
> > do you think a neural network, be it natural or artificial,
> > accomplishes? What does it do? How is its operation correlated to the
> > environment, say, during simple categorization (of audio or images)?
>
> It filters and organises input data. That is, given a mess of inputs,
> the NN passes some and blocks others. The output may be "pattern
> recognition", or input to another NN. The following description will
> show what I understand happens: Given an NN that's been trained to
> "recognise the letter A". Input visual data of the letter B. Some of
> that data will "pass", that is, activate some connections as the data
> flows through the NN (eg, the centre horizontal bar of the B.) Most will
> not. The NN will respond (output "Letter A") if enough of the
> connections within it are activated and not otherwise. The "reference
> state" is the organisation of the connections. Eg, a certain group
> inputs are connected to one element, some of the same inputs are
> connected to another, and some to a third. If all inputs are on, all
> three elements will fire; otherwise only one, or none. NB that my
> description assumes a layered architecture. A good example of what I
> envision is the mammalian retina, which is a neural network of exquisite
> precision. IMO the organisation of the retinal NN is repeated throughout
> the brain.
>

This description is basically ok, except you have to overlay it on the anatomy of the cortex.

In an artificial NN, all you need is one layer of cells − ie, simple perceptrons − to do the letter recognition. and in which case, the important info **is** stored in the weights of those cells only. In contrast, as described above, it seems to me much more likely that the real brain works with a more layered structure. The lower 3 layers in the retina, plus another 6 or 8 in the cortex, are mainly involved in filtering and not in "storage". IOW, they filter and perform the basic component operations which occur prior to eventual recognition. They don't store the letter "A", for instance, they filter the top and side edges of the "A", filter the color, filter the a blob exists in the region of visual space of the "A", on and on. Then all of these signals are compared with each other [via those 1100+ interconnection pathways − that just won't go away], and also compared to the actual "stored" signals in the association cortex. Those signals were stored previously, as a result of this entire same process having occurred previously on signals which were ultimately deemed to be "meaningful".

This last is related to general emotional or motivational inputs from brain stem and amygdala and hypothalamus essentially directing the association areas regarding whether or not to permanently store, or disregard, the outputs for the visual filters. This is a nice simple formulation [albeit the implementation is somewhat complex], and gets away from the problems you've been expressing above.
=================

> *The "correlation to the environment" is a phrase that I don't find very*
> *meaningful. There are organised inputs from the sensors. If the sensors*
> *provide inputs such that the "A−recognition NN" responds, then the*
> *system has "seen" the letter A.*
>

Yeah, after you go thru all the many steps described above. Vision *is* special, and that's why 1/3 or so of the cortex and 30+ visual areas are devoted to it. When I read your descriptions, I keep seeing you arguing a 2−level system. It's not.
===================

> *And categorisation of audio or images is _not_ simple. Anyone who has*
> *tried to write a categorisation tree to distinguish between CAT and DOG*
> *should know that. I've doodled some myself, and I've found that at every*
> *level there are assumptions about what the subcategory "means", ie, each*
> *node hides another tree. Add to that the problem of recognising a CAT*
> *or DOG from variable visual data, and the problem is seen to be*
> *horrendously complex. It's no accident that optical NNs have barely been*
> *ablt to recognise letters, and even those not reliably. IMO, the central*
> *problem of NNs is not their general organisation or properties, it's the*
> *nasty little details of how to actually connect the actual elements.*

Yes, it's that the [artificial] systems currently in use aren't nearly complex enuf to be able to handle the problem.

It took evilution 3.5BY to work it out. If you will recall, we discussed previously that vision first appeared early on during the Cambrian explosion in trilobites 550MYA or so. As soon as multicellular creatures started to flourish, vision evolved. And it's been evolving ever since. The original creatures saw little more than light and dark, and possibly movement, including even the much later−evolved frogs, as mentioned earlier. But frogs don't cross busy highways very well, and it took evolution of all of those 30+ visual areas, each with their more or less specific processing jobs, before we really had something truly and generally useful.
===================

> *Small differences in connection topology may have large effects on the*
> *NNs operation; and those effects aren't easily predictable. It's much*
> *easier to refer to "emergent properties" than to design a device that*

> *will display those emergent properties.*
>
> *Corrections and clarifications welcome.*

Yes.