

Re: The Hard Problem for Behaviorists

Source: <http://sci.tech-archive.net/Archive/sci.cognitive/2004-11/0078.html>

From: Curt Welch (curt_at_kcwc.com)

Date: 11/02/04

Date: 02 Nov 2004 20:07:50 GMT

"JPL Verhey" <matterDELMinds@hotmail.com> wrote:

> *"Curt Welch" <curt@kcwc.com> wrote in message*

> *So in way talking about process and structure is actually a description*
> *of the interpreting brain itself. We are hard-wired/programmed to see*
> *everything we see in those terms. Seems relevant to AI?*

Well, I don't think the ideas of "processes" and "structure" are the correct low level abstractions to define the operation of the brain with – as I went on to talk about below. But, yes, in our search to understand how the brain does what it does, we must identify the correct low level primitives to describe it's operation. Until we do so, we will probably not be able to understand how it works. Maybe we already have the correct low level primitives and concepts to work with, or maybe we don't.

> > *I think trying to understand the universe in terms of "structure" and*
> > *"processes" is just not the way to do it. And I don't think that's*
> > *how the*
> > *brain works either.*
>
> *But to see structure and process, or text and context, tells in any case*
> *something of the "output" of our interpreting brain – our experiences.*
> *To understand better how this comes to be, we have to back-track in the*
> *opposite direction of the causal arrow the events that give rise, in*
> *the end, to our experiences.*

Yes, that's true.

What's important to me is that we have been trained how to talk about our experience. We learned how to do it. It was not hard wired into us by our genetics. If we choose to use terms like "structure and process", and to "see structure and process" in our experiences, and to think about our experience as structure and process, then that is not an indication of how the brain hardware works, that is only an indication of what our brain was trained to do by the environment it was raised in.

> > *The language we use to describe how computers operate does have close*

> > *parallels to these concepts of structure and processes. We think of*
> > *computers as state machines, which transform from one well defined*
> > *state,*
> > *to the next, in a predictable sequence. So each state the machine is*
> > *in*
> > *becomes a "structure", and the progression from one state to the next*
> > *is*
> > *our description of the "processes".*
>
> *But how does this transition from one state to another (the process)*
> *actually occur?*

I don't understand why you ask this. Do you not know how computers work?
Or are you trying to make some point which I do not see?

State machines have an external power source which keeps them moving. The structure of the machine guides the motion from one state to the next. The external power source is what makes it change states, and the structure of the machine is what controls the sequence of states which get visited.

A wheel is an example of one of these machines. If you attach it to an external power source, you can make it spin. If you put a mark on the wheel, then you can define states based on the position of the wheel. You can define four states based on points of the compass. Whichever point the mark is closest to is the state of this machine. As the wheel spins, the machine changes from state N, to E, to S, to W and back to N. Or if it spins in the other direction, the state change is reversed.

The structure of this machine defines the states, and controls the possible state transitions. It's impossible for this machine to change from state N, to state S, without visiting either E or W first.

So the external power source is what makes it change state, and the physical structure of the machine is what controls which states it changes to.

Computers work the same way, they just have a lot more states, and our modern electronic computers are pulsing electrons around instead of wheels.

> *But aren't those "value's" attached by us to input signals, that cause*
> *computer-states to follow each other sequentially? So the value predicts*
> *a future state of the computer. However the "state" of the computer as*
> *an abstraction created in by the human brain. i.e. how we observe the*
> *computer and think of it.*

I assigned the "state names" to the position of the wheel in the example above. In doing so, I created a language to talk about the operation of the wheel. So then later I could say the wheel changed from state N to state E. But, behind my language, is a wheel, doing what it does, no matter how I choose to talk about it. I could have for example created 360 "states" instead of 4, and called them 0 to 359, and talk about how the wheel

changed from state 0 to state 90, passing through all the numbers between there as it did it. My language is different at that point, but the operation of the machine has not changed.

So yes, our description can vary a lot yet the machine still does "the same thing". So to describe how a brain works, we don't have to find the one and only "correct" language to use to describe it. There are no doubt many ways to describe it and still understand how it works. But if you pick the wrong language, the description becomes much more complex – and many times, so complex as to be impossible to understand. Describing the orbit of the planets around the earth is just as valid as describing the orbit of the planets around the sun. But when you do it in terms of the sun, the orbits become very simple to explain (requires few words), but when you do it placing the earth at the center, the description is so complex that it's almost impossible to understand. A small change in your language can make a huge difference in the complexity (and typically also the accuracy) of your description. This is why it's so important to find the correct language to use in order to describe the operation of the brain.

> > *But then I saw the real beauty behind pulse signals. They have no
> > structure (or state), because the input value is never "constant" for
> > some
> > period of time. Each pulse is simply a temporal event marker. There
> > is no
> > constant state represented in the data when you use temporal events as
> > the
> > low level basis for the operation of the machine. The network/machine
> > is a
> > dynamic event processor, and has no "states".
>
> *I would think that actually and independent of the interpreting
> observer, even a binary "solid state" computer is never a constant.**

No, of course not. Just like my wheel machine example above, the parts of a computer are in constant fluid motion when it's running. It never stands still.

> > *This type of machine operates by dynamically classifying events. So,
> > instead of creating an "output state" as a function of the "input
> > state +
> > machine state", this type of machine simply maps input events to
> > output
> > events.
>
> *What does the last sentence "maps input events to output events" mean?**

Well, this is all the language I use to talk about this type of machine.

As a mechanical analogy, you can think of this type of network like you would a back box with X holes in the top, and Y holes in the bottom. You

drop marbles in the top, and every marble you drop in comes out one of the bottom holes. The black box simply routes marbles to the correct output hole. It can route it to different output holes at different times. So when you drop a marble in hole X1, it may come out Y4. But the next marble you drop in X1 might come out Y10. It's a temporal network so the timing, and not just the sequence, is important in determining where the marbles come out.

Each marble is what I talk about as being an "event". The box creates a mapping function from input holes to output holes.

Each marble, or "event", is the same as a pulse in a neural network. But in my network, events are conserved, so one marble in always creates one marble out. You never have the pulse die in the network, or trigger multiple pulses to come out. The conservation property converts the network into a strict mapping function, from one event in, to one event out.

BTW, it's a reinforcement learning network, so my black box would need extra button on it to allow it to be punished and/or rewarded.

> > *Constant "structure" or "state" is represented in this language by a*
> > *regular flow the same type of event.*
>
> *But the flow/behavior of electricity in such a machine depends on*
> *certain threshold values?*

I have no idea what you are getting at when you ask that. Electrical machines do tend to function differently when you exceed their design limits, yes. And a pulse processing network would have threshold levels which separated a pulse from a non pulse state on a wire. So when you send the network a pulse, you would have to be the correct amount of electricity or else the machine would not work as expected.

And if I built this network out of marbles and mechanical parts (which I could do), then you would have to put the correct sized marbles in the input holes or else the machine would not work as expected.

> > *If this is how our brain works and how we understand everything at the*
> > *lowest level, then to us, the world is not "structure" and*
> > *"processes",*
> > *it's just a constant flow of temporal events. It never static. It's*
> > *100%*
> > *dynamic.*
>
> *Yes.. and that never static world includes binary solid state computers!*

Ah, so maybe I get where you have been going with your comments. I'm not trying to argue that computers can't be intelligent because computers don't work correctly. I can implement my network on a normal digital computer with no problem at all. I'm just saying that the language we normally use

to design computer hardware and software tends to lead us in the wrong direction. It causes us to think about the problem incorrectly. State machines move through their state space as quickly as possible. The state they change to next is normally a function of only the last state, and not a function of time. My network is a temporal network, and the state it changes to is totally a function of how long it's been in the current state.

You can't even specify the operation of this type of network with a Turing machine for example. But because real computers do have real time clocks which are part of their state, I can easily program the time dependent function into a normal digital computer.

You can create a simulation of my type of network on Turing machine for example, but without the real-time clock input, you can't connect it to the real world.

> > > *and everything I understand about the physical*
> > > *world comes to me through the senses.*
> > >
> > > *It's as if we were locked inside a control room in our heads*
> > > *looking*
> > > *at the*
> > > *world through video screens and controlling our body with joy*
> > > *sticks.*
> > > *It*
> > > *feels to us as if we were actually seeing and touching the real*
> > > *world,*
> > > *but*
> > > *in fact, we can't do that, because "we" don't exist in the real*
> > > *world.*
> > > *We*
> > > *are forced to exist in this private virtual world.*
> > >
> > > *Yes, that's a very good analogy. But I would not call it a "virtual*
> > > *world" at all! It is, in fact, the most real world thinkable.*
> >
> > *Yes, I agree. However, think about what that means. It means that*
> > *what we*
> > *think of as "virtual" is just as real as what we think of as "real".*
>
> *I think it means that in any case we conclude there are two different*
> *entities – one the representation, and the other the represented. To*
> **us* the representation is the real one – ie we *are* that*
> *representation, there is no "observer in representation" – whereas the*
> *represented is hidden. But there also is the bigger reality to which*
> *both the represented and the representation belong.*

Sure. With the exception that it's still unclear if the two really exist in the same reality. I think they do, but it's not yet proven.

sci.cognitive: Re: The Hard Problem for Behaviorists

> > *If the "real universe" to us is nothing more than the signals in our
> > head,
> > then why should we look at the signals in a computer which define a
> > "virtual" world in a video game, as being any less real?*
>
> *The representation, i.e. the flow of our experiences IS the real world,
> for each of us. Those are not "signals in our heads". "Signals in our
> heads", brains are also visual representations, thoughts, abstractions.
> There is nothing virtual about them, nor defining a video game and run
> it on a computer. There is even nothing virtual about calling the video
> game when you play it "a virtual reality". Or call it dream game! Or
> just a game fun to play.. There is nothing not-real about those
> descriptions.*

Right, I think we are saying the same thing. You just choose to replace the word "virtual" with "real", and I chose to replace "real" with "virtual". What we both seem to be saying is that "real" and "virtual" are the same thing.

--

Curt Welch
curt@kcwc.com

<http://CurtWelch.Com/>
<http://NewsReader.Com/>