

## Re: AI and cognitive psychology rant (getting more and more OT – tell me if I should shut up)

*Source:* <http://coding.derkeiler.com/Archive/Python/comp.lang.python/2003-10/2218.html>

---

*From:* Stephen Horne (\$\$\$\$\$\$\$\$\$\$\$\$\$\$\$\$\_at\_\$.co.uk)

*Date:* 10/15/03

Date: Wed, 15 Oct 2003 10:50:42 +0100

On Tue, 14 Oct 2003 16:58:45 GMT, Alex Martelli <aleax@aleax.it> wrote:

>Isaac To wrote:

>

>>>>>> "Alex" == Alex Martelli <aleax@aleax.it> writes:

>>

>> Alex> Chess playing machines such as Deep Blue, bridge playing  
>> programs Alex> such as GIB and Jack ..., dictation-taking programs  
>> such as those Alex> made by IBM and Dragon Systems in the '80s (I  
>> don't know if the Alex> technology has changed drastically since I  
>> left the field then, Alex> though I doubt it), are based on  
>> brute-force techniques, and their Alex> excellent performance comes  
>> strictly from processing power.

>>

>> Nearly no program would rely only on non-brute-force techniques. On the

>

>Yes, the temptation to be clever and cut corners IS always there...;-).

It's an essential part of an intelligent (rather than perfect) solution.

The only one of the programs above where I know the approach taken is Deep Blue. A perfect search solution is still not viable for chess, and unlikely to be for some time. Deep Blue therefore made extensive use of heuristics to make the search viable – much the same as any chess program.

A heuristic is by definition fallible, but without heuristics the search cannot be viable at all.

So far as I can see, this is essentially the approach that human chess players take – a mixture of learned heuristics and figuring out the consequences of their actions several moves ahead (ie searching). One difference is that the balance in humans is much more towards heuristics. Another is that human heuristics tend to be much more

sophisticated (and typically impossible to state precisely in words) because of the way the brain works. Finally, there is intuition – but as I believe that intuition is basically the product of unconscious learned (and in some contexts innate) heuristics and procedural memory applied to information processing, it really boils down to the same thing – a heuristic search.

The human and computer approaches (basically their sets of heuristics) differ enough that the computer may do things that a human would consider stupid, yet the computer still beat the human grandmaster overall.

Incidentally, I learned more about heuristics from psychology (particularly social psychology) than I ever did from AI. You see, like Deep Blue, we have quite a collection of innate, hard-wired heuristics. Or at least you lot do – I lost a lot of mine by being born with Asperger syndrome.

No, it wasn't general ability that I lost by the neural damage that causes Asperger syndrome. I still have the ability to learn. My IQ is over 100. But as PET and MRI studies of people with Asperger syndrome reveal, my social reasoning has to be handled by the general intelligence area of my brain instead of the specialist social intelligence area that neurotypical people use. What makes the social intelligence area distinct in neurotypicals? To me, the only answer that makes sense is basically that innate social heuristics are held in the social intelligence area.

If the general intelligence region can still take over the job of social intelligence when needed, the basic 'algorithms' can't be different. But domain-specific heuristics aren't so fundamental that you can't make do without them (as long as you can learn some new ones, presumably).

What does a neural network do under training? Basically it searches for an approximation of the needed function – a heuristic. Even the search process itself uses heuristics – "a better approximation may be found by combining two existing good solutions" for recombination in genetic algorithms, for instance.

Now consider the "Baldwin Effect", described in Steven Pinkers "How the mind works"...

.....

So evolution can guide learning in neural networks. Surprisingly, learning can guide evolution as well. Remember Darwin's discussion of "the incipient stages of useful structures" – the what-good-is-half-an-eye problem. The neural-network theorists Geoffrey Hinton and Steven Nowlan invented a fiendish example. Imagine an animal controlled by a neural network with twenty connections, each either excitatory (on) or neutral (off). But the network is utterly

useless unless all twenty connections are correctly set. Not only is it no good to have half a network; it is no good to have ninety–five percent of one. In a population of animals whose connections are determined by random mutation, a fitter mutant, with all the right connections, arises only about once every million (220) genetically distinct organisms. Worse, the advantage is immediately lost if the animal reproduces sexually, because after having finally found the magic combination of weights, it swaps half of them away. In simulations of this scenario, no adapted network ever evolved.

But now consider a population of animals whose connections can come in three forms: innately on, innately off, or settable to on or off by learning. Mutations determine which of the three possibilities (on, off, learnable) a given connection has at the animal's birth. In an average animal in these simulations, about half the connections are learnable, the other half on or off. Learning works like this. Each animal, as it lives its life, tries out settings for the learnable connections at random until it hits upon the magic combination. In real life this might be figuring out how to catch prey or crack a nut; whatever it is, the animal senses its good fortune and retains those settings, ceasing the trial and error. From then on it enjoys a higher rate of reproduction. The earlier in life the animal acquires the right settings, the longer it will have to reproduce at the higher rate.

Now with these evolving learners, or learning evolvers, there is an advantage to having less than one hundred percent of the correct network. Take all the animals with ten innate connections. About one in a thousand (210) will have all ten correct. (Remember that only one in a million nonlearning animals had all twenty of its innate connections correct.) That well–endowed animal will have some probability of attaining the completely correct network by learning the other ten connections; if it has a thousand occasions to learn, success is fairly likely. The successful animal will reproduce earlier, hence more often. And among its descendants, there are advantages to mutations that make more and more of the connections innately correct, because with more good connections to begin with, it takes less time to learn the rest, and the chances of going through life without having learned them get smaller. In Hinton and Nowlan's simulations, the networks thus evolved more and more innate connections. The connections never became completely innate, however. As more and more of the connections were fixed, the selection pressure to fix the remaining ones tapered off, because with only a few connections to learn, every organism was guaranteed to learn them quickly. Learning leads to the evolution of innateness, but not complete innateness.

""""

A human brain is not so simple, but what this says to me is (1) that anything that allows a person to learn important stuff (without damaging flexibility) earlier in life should become innate, at least

to a degree, and (2) that learning should work together with innateness – there is no hard divide (some aspects of neurological development suggest this too). So I would expect some fairly fixed heuristics (or partial heuristics) to be hard wired, and I figure autism and Asperger syndrome are fairly big clues as to what is innate. Stuff related to nonverbal communication such as body language, for instance, and a tendency to play activities that teach social stuff in childhood.

And given that these connectionist heuristics cannot be stated in rational terms, they must basically form subconscious processes which generate intuitions. Meaning that the subconscious provides the most likely solutions to a problem, with conscious rational thought only handling the last 5% ish of figuring it out. Which again fits my experience – when I'm sitting there looking stupid and people say 'but obviously it's z' and yes, it obviously is – but how did they rule out a, b, c, d, e, f and so on and so forth so damned quick? Maybe they didn't. Maybe their subconscious heuristics suggested a few 'obvious' answers to consider, so they never had to think about a, b, c...

Human intelligence is IMO not so hugely different from Deep Blue, at least in principle.

*>no understanding, no semantic modeling.*

*>no concepts, no abstractions.*

Sounds a bit like intuition to me. Of course it would be nice if computers could invent a rationalisation, the way that human brains do.

What? I hear you say...

When people have had their corpus callosum cut, so the left and right cerebral cortex cannot communicate directly, you can present an instruction to one eye ('go get a coffee' for instance) and the instruction will be obeyed. But ask why to the other eye, and you get an excuse (such as 'I'm thirsty') with no indication of any awareness that the request was made.

You can even watch action potentials in the brain and predict peoples actions from them. In fact, if you make a decision in less than a couple of seconds, you probably never thought about it at all – whatever you may remember. The rationalisation was made up after the fact, and 'backdated' in your memory.

Hmmm – how long does it take to reply to someone in a typical conversation...

Actually, it isn't quite so bad – the upper end of this is about 3 seconds IIRC, but the absolute lower end is something like half a second. You may still have put some thought into it in advance down to

comp.lang.python: Re: AI and cognitive psychology rant (getting more and more OT – tell me if I should shut up)

that timescale (though not much, obviously).

Anyway, why?

Well, the textbooks don't tend to explain why, but IMO there is a fairly obvious explanation. We need excuses to give to other people to explain our actions. That's easiest if we believe the excuses ourselves. But if the actions are suggested by subconscious 'intuition' processes, odds are there simply isn't a rational line of reasoning that can be put into words. Oh dear – better make one up then!

*>Again, that't not AAAI's claim, as I read their page. If a problem  
>can be solved by brute force, it may still be interesting to "model  
>the thought processes" of humans solving it and implement \_those\_ in  
>a computer program -- \_that\_ program would be doing AI, even for a  
>problem not intrinsically "requiring" it -- so, AI is not about what  
>problem is being solved, but (according to the AAAI as I read them)  
>also involves considerations about the approach.*

1. This suggests that the only human intelligence is human intelligence. A very anthropocentric viewpoint.
2. Read some cognitive neuroscience, some social psychology, basically whatever you can get your hands on that has cognitive leanings (decent textbooks – not just pop psychology) – and then tell me that the human mind doesn't use what you call brute force methods.

--

Steve Horne  
steve at ninereeds dot fsnet dot co dot uk