

From: Joscha Bach <[REDACTED]>
To: "S.M. Kosslyn" <[REDACTED]>
Cc: Jeffrey Epstein <[REDACTED]>
Subject: Re: Today's discussion
Date: Tue, 19 Feb 2013 03:16:16 +0000

Dear Stephen,

Berlinale is over and I got back to your latest message, and the matter of IQ tests for AI in particular.

>> But nonverbal thinking is something that I suspect is quite similarly powerful in other primates.

>

> I think we are much better at this than other primates; our conceptual structures are more powerful, and they in turn drive more powerful mental simulations

I remember a very vivid discussion at a conference for cognitive modeling (John Anderson and a few others were on the panel), wrt the cognitive differences between humans and apes. No-one seemed to agree on anything; it was beautiful. BTW, John suggested at the time that it would be mainly quantitative and not qualitative. Chimps would theoretically be able to learn everything a human could, but not in a single life-time.

From teaching computer science, I got the impression that there are conceptual structures that cannot be taught to all humans (at normal IQ levels). There is some weak research to support this: the ability to understand the concept of variables and pointers seemed to be quite invariant regardless of whether the students had taken only two hours or a complete course. The argument is weak, because it is difficult to exclude the possibility that a completely different didactic trajectory would have resulted in success.

However, it seems clear that teaching programming requires grasping a variety of principles that provide students with different challenge levels (variables, pointers, iteration, recursion, closures and currying, higher order functions, etc.). If it turns out that certain programming techniques cannot be taught to average humans (in the sense of a general IQ of around 100), it would mean that humans are not generally intelligent. But I guess that would not come as a big surprise.

> Grammar is no doubt important, but I'm just not sure that it's at the root of what's most interesting about human intelligence.

What do you have in mind?

>>> (...) The WAIS has some 11 subtests, which cover a wide range of underlying abilities (and are much more challenging)

>>

>> Lets look at them (I have to admit that I am no expert on this, and it is quite some time ago that I looked at IQ testing):

>> - The processing speed tests are probably trivial for computers

>

> If memory serves, none of the tests are about processing speed per se -- they are timed, but the issue is not simple processing speed, its facility with certain kinds of reasoning

Yes, but humans generally cannot brute force these tests. Quite often, it amounts to a similar thing as it did with chess: the Shannon A vs. the Shannon B strategy. Roughly put, Shannon suggested that there are two ways to play chess. One would be to come up with a devious, long-winded plan, and try to proof it against mistakes and failure. The other one would be to exhaustively search through the space of possible game trajectories, while

discarding unpromising dead-ends as early as possible. The former strategy is used by human players, and made better by subtly trained pattern matching over time, and the latter is played by computers up to this day. For many tests, we can replace reasoning facility (the subtle trained patterns) with sheer computational power and perfect memory.

>> - The working memory tests are likewise rather simple engineering problems

>

> Again, none of the tests specifically assess WM, although several tap into it.

I was not clear, sorry. If I understood correctly, the WM problem set contains mental arithmetic and number sequence learning. This is what would make them exceptionally simple for computers, in my opinion. (If they were really assessing the structure, abilities and connectivity of WM, we'd be in trouble.)

>> - Perceptual reasoning is somewhat similar to the Raven (maybe I underestimate them?)

>

> There are a set of perceptual reasoning tests, only some of which are at all like Raven

>

>> - Verbal comprehension:

>> - similarities and vocabulary tests are classical AI and computational linguistics

>> - information is close to IBM's Watson (recognition and inference)

>

> SO.. what you seem to be saying is that it would be simple to program a computer to do well on IQ tests. I would love to see this!

AI has been diddling with human IQ tests since the 60ies, but the idea crops up every now and then. (I have heard it from Selmer Bringsjord, who thinks that we need new, improved IQ tests for that.)

There seems to be somewhat of a consensus that traditional IQ tests can usually be solved quite well by a narrow AI solution. A quick Google search comes back with a recent attempt by Claes Strannegård. On the other hand, he only attacked (and bested) the Raven.

> I think most of the above is in fact implicit in some of the tests. Remember that factor analysis reveals a very rich structure of human intelligence, with 60+ specific identifiable abilities that feed into it

Certainly! But if humans used over 60 tools to expertly chisel a block of marble, it could still mean that a CNC machine can do it with only three! Also, humans are the only animals that are capable of turning marble into portrait busts, and yet that ability would not be a good benchmark for artistic prowess in a CNC machine.

That being said I believe that we should administer IQ tests to our cognitive architectures. I only suspect that a test solving machine might not be generally intelligent, because it might do so without the basic qualitative functionality that we take for granted in all test subjects, and that so far eludes us.

>> Please tell me if my take on the WAIS is wrong!

>>

> I think you might enjoy actually taking it. (My wife, when she was in training, used me as a guinea pig for testing -- and I found taking the test really interesting...)

> (...)

> Better yet: Have somebody actually give it to you. The actual WAIS cannot be taken on a computer or the like; it needs a trained person to administer it

You make it so intriguing; it sounds so much like LSD or Psylocibine ;-)

> It may be reflected, but such discourse is not a necessary consequence of intelligence. A deaf mute could still

be very intelligent.

I assume that you were not just referring to spoken words, but also include sign language, joint construction of artifacts etc. which also can be construed as discourse, and thus our example deaf mute would also be afflicted with a diminished sense of touch and sight?

Interesting distinction that you open up here: intelligence as a functional potential vs. a realizable performance. While we would never say about a paraplegic man that he could still be very fast, intelligence as potential makes much more sense, since it reflects an ability to process information, even in the absence of absorbing and dispensing said information at a high rate.

However, in the context of the above argument, we could supply our deaf mute AI with perfect communicative prostheses. I guess that if it is generally intelligent, it should be able to make good use of them, and if it is not, then it won't be able to fake their use.

>> which includes interpreting and creatively structuring the world. Many of the things that the WAIS measures, like recognizing and categorizing shapes, are prerequisites for that. Others might be acquired tastes that emerge on more basic functionality, like mental arithmetic. But a toolbox is not an architecture. A collection of tubes, tires, pedals and spokes is not a bicycle.

>

> Good distinction. The IQ tests require a suite of skills and abilities, which could in principle arise from numerous architectures..

Oh, so we seem to be in agreement. (And yet I will have to look at the WAIS in much more detail.)

> You must be familiar with what the classic AI guys (e.g., Herb Simon) called "the representation problem"

Jeffrey just asked me to re-visit Roger Schank's SPGU book. His exclusive focus on episodic memory (among many other things) would be inadequate today, but overall, it has aged well over the last 35 years. At the the beginning of his conceptual dependency theory, he points out:

1) For any two sentences that are identical in meaning, regardless of language, there should be only one representation.

2) Any information in a sentence that is implicit must be made explicit in the representation of meaning of that sentence.

The first one is clear (especially if we ignore activation patterns, which would affect associations, subsequent use of pronouns etc.), and IMHO means that the classical cognitive architectures won't cut it: there are infinitely many semantically equivalent symbolic representations. (I think we solve this by ensuring that the architecture constrains the space of possible representations down to one.)

Number two is a little controversial in the light of the extended mind, because some of the implicit information can reside in the world or in systemic affordances that are not explicitly represented. But apart from that bit, I think that Roger Schank had that one right, too, and it is where the classical linguists (and especially Jerry Fodor right up until today) went wrong: meaning is not outside of the representations, at the receiving end of some metaphysical arrow (or rather, if it is, then the notion of "meaning" itself is entirely misconstrued and broken, and should be replaced by something more innocent, like "encoding".)

To operate efficiently with representations, they need to be organized with respect to both the available operations (like constrained activation spreading, planning, linguistic reference, reflection/internal perception etc.) and the things we store them for, i.e. motivational relevance. I think that especially motivational relevance might be key for overcoming the representation problem in cognitive model with limited power.

However, I might be wrong at least in the case of low level sensory processing. Last year, Andrew Ng and a bunch of Googlers extracted object categories from 30 million random Youtube frames, using an unsupervised

neural network. Mere statistical structure, when put into a suitably designed computational sieve, might already suffice to filter out many of the most interesting regularities.

> I would stop before language, but this may reflect a deep prejudice on my part. I think that much of logic comes out of perceptual experience with contingencies in the world

That is no contradiction. I am sure that chimps also possess much of logic, and yet they won't be able to learn how to program.

>> emotion is highly interesting, that Damasio is quite correct with respect to what emotion does, and that it makes a lot of sense (and is fun) to equip AIs with emotion, mood, affect and emotional dispositions. But strictly necessary? No.

>

> I disagree; I think emotion is crucial for rapid interrupts and setting priorities (yes, motivation is also involved, but generally has a longer time horizon)

That is one of the roles that it plays in humans, but in an artificial system, we can probably overcome the need for rapid interrupts and resource allocation by adding more hardware. Instead of switching priorities, we might use multiple systems in parallel, and instead of a startle mechanism, we might use an additional monitoring CPU. What it boils down to is: emotion reduces the complexity of cognitive processing operations by tailoring them to the given task and environmental dynamics. If that complexity reduction results in some factor of n for processing time and/or resource use, then we can overcome this with an investment into more hardware or better algorithms. If the reduction is exponential or better, then emotion might be necessary in a strict sense. Ultimately, the question is an empirical one, but my bet is on the former option.

The message has gotten quite long again; sorry for that. I find that it very helpful to have an opportunity to spell these arguments out! Thank you, and cheers!

Joscha