
From: Joscha Bach <[REDACTED]>
Sent: Thursday, February 14, 2013 11:42 PM
To: S.M. Kosslyn
Cc: Jeffrey Epstein
Subject: Re: Today's discussion

Dear Stephen,

thank you for your so far unending patience in that discussion. --> <http://www.xkcd.com/386/>

>> I agree, and yet only a subset of animals can use symbols for =ommunication.
>
> communication and intelligence are not the same thing (think about =ees etc)
>
>> Of these, only a small subset can make use of negation, conjunctions =nd disjunctions in symbolic communication (for instance, Irene =epperberg's famous grey parrots). It appears that only humans can learn =ich grammatical language, and I suspect that this is the primary =nabler of our superior problem solving capabilities.
>
> I disagree. Einstein claimed that his greatest discoveries came from =ental imagery, and he later converted those thoughts to verbal =xpressions only with great difficulty. I think AI has vastly =nderestimated the role of "mental simulation/emulation" in thinking and =easoning.

Nonverbal thinking is primary, and is poorly understood, and it has been =eglected by what we now often call "classic AI". (But if we look at the =original ideas, that was probably not intentional. Logic based systems =ere low hanging fruit. Once you have a paradigm and a community, you =nd up with a methodology that is bound to stay, unfortunately.)

But nonverbal thinking is something that I suspect is quite similarly =owerful in other primates. I think that the most interesting difference =etween chimps and humans is how we can use grammatical language to =program" and organize our thought processes, and how we can use it to =uggest, transmit, create and manipulate new categories.

I am on your side insofar as I think that the important research needs =o be done in mental imagery (or more accurately: on mental =epresentations and operations that facilitate mental imagery, among =ther things). But I think that human intelligence is shaped by the =dditions of grammars, which happen to be relatively easy to implement =hen you look at them in isolation. Only grammar on its own cannot do =hat Einstein did.

>>
>>> --]]]]]]]]]]]]]]]]]] So.. what would be wrong with building a machine =hat could do well on IQ tests?
>>
>> Nothing is wrong with building a machine that excels at playing chess =r cooking coffee or scoring that the Raven test.
>
> The Raven would be a bad idea -- way too easy. The WAIS has some 11 =ubtests, which cover a wide range of underlying abilities (and are much =ore challenging)

Lets look at them (I have to admit that I am no expert on this, and it =s quite some time ago that I looked at IQ testing):
- The processing speed tests are probably trivial for computers
- The working memory tests are likewise rather simple engineering =roblems
- Perceptual reasoning is somewhat similar to the Raven (maybe I =nderestimate them?)
- Verbal comprehension:

- similarities and vocabulary tests are classical AI and computational linguistics
- information is close to IBM's Watson (recognition and inference)

The only thing that looks interesting to me in the WAIS is the comprehension test, because I don't see a straightforward approach to heat on them with narrow AI. I would like to expand exactly this domain: making sense of the world.

We don't have advanced problem solving ("these are the rules for chess. How would you try to beat a beginner level player, a medium player, a top player most quickly?"). We don't have constructive abilities. We don't have verbal creativity etc.

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

> How do you know for sure what the "basics" are?

While the literal understanding of the Turing Test leads nowhere (or, well, to the Loebner prize), I think that he had the right idea. Intelligence is reflected in the ability to participate in meaningful discourse, 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, wires, pedals and spokes is not a bicycle.

Some of the basics stem directly from the requirements of producing adequate representations of perceptual and abstracted content (hybrid hierarchical representations that can do associations, compositional stuff, grammatical systematicity, learning and categorization, inheritance etc.). Others come from the needs to get the processes of bottom-up/top-down perception, reflection, memory retrieval, inference, analogy building etc. to work. And some have to do with the requirements of translating between Mentalese (in Pinker's sense, not in Fodor's) and natural language.

I may delude myself in thinking that I know what the basics are. In fact, it is extremely likely that I do (every computer science problem seems to be misconceptualized until it has been properly implemented). But I would start with mental representation, perceptual processing and motivational relevance, and then go for language, while revisiting those areas that turn out to fall short.

> Forget about the Raven; it's a non-verbal test of fluid intelligence which in fact turns out to have, by accident not design, two different types of items -- solved by spatial vs. analytic strategies). The Raven does not even begin to characterize all of what is captured by the WAIS

It might well be that I totally underestimate the WAIS requirements; I will look at them.

>> I agree. But I am not convinced that proper emotions are absolutely necessary for Intelligence (motivation might suffice to drive some kind of non-emotional, serene Buddha intelligence). I am nevertheless interested in understanding and modeling them.

>

> I think Antonio Damasio and his successors have made a very good case that emotion plays a key role in reasoning. (Not just motivation, actual motion.)

Most of what I would say against that has been better expressed by Aaron Sloman. For instance, if my computer is prone to swapping memory content to hard drive and back, and I kill the part of the OS that coordinates the swapping, my computer is likely to malfunction. But from this I cannot infer that computers cannot work without swapping. Damasio's argument does not convince me because he does not elucidate a functional role that would emotion an absolute requirement for an artificial mind. Lesion studies amount to shutting down parts of an operating system that has been designed to cope with very specific requirements. I believe that in humans, emotions structure social interaction, support communication, prime memory and cognitive processing, and most importantly, allocate the scarce resources of our mind according to the current situation. None of this is necessary if I remove the resource constraints.

But again, perhaps there are better arguments now than in the original =omatic marker hypothesis?

Please do not misunderstand me; I am making a merely philosophical point =ere, with respect to the basic requirements for Intelligence. I think =hat emotion is highly interesting, that Damasio is quite correct with =espect to what emotion does, and that it makes a lot of sense (and is =un) to equip AIs with emotion, mood, affect and emotional dispositions. =ut strictly necessary? No.

> Are you involved in BICA? That seems like a natural community for you!

The way I understand it, there are at least four very similar groups =ow: cognitive modeling (that is where John Anderson goes), AGI (started =y Ben Goertzel as an attempt to revive the original AI), BICA (a =emnant from the failed DARPA proposal of the same name, and later =icked up by Alexei Samsonovich as an alternative to AGI, I suspect =ecause he does not get along with Ben), and Cognitive Systems (Pat =angley et al.). I basically like them all, and think that they should =oin forces, while simultaneously raising the bars against narrow AI and =cience fiction. Many members of the audience already belong to two or =even three of the groups. Alas, politics, mutual accusations of =cruffiness and stuffiness, and so on...

Personally, I have not been to one of the BICA conferences (only a =ouple planning workshops), and I am on their roster of reviewers.

Cheers,

Joscha

```
<?xml version=.0" encoding=TF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN" "http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version=.0">
<dict>
  <key>conversation-id</key>
  <integer>245038</integer>
  <key>date-last-viewed</key>
  <integer>0</integer>
  <key>date-received</key>
  <integer>1360885332</integer>
  <key>flags</key>
  <integer>8623750145</integer>
  <key>gmail-label-ids</key>
  <array>
    <integer>6</integer>
    <integer>2</integer>
  </array>
  <key>remote-id</key>
  <string>276458</string>
</dict>
</plist>
```