
From: S.M. Kosslyn [REDACTED]
Sent: Friday, February 15, 2013 4:24 AM
To: Joscha Bach
Cc: Jeffrey Epstein
Subject: Re: Today's discussion

Hi J..

Some responses and reflections below...

||||||||||||||||||||||||||||||||||||||||

Stephen M. Kosslyn
Director, Center for Advanced Study in the Behavioral Sciences Stanford University
75 Alta Road
Stanford, CA 94305
Voice: [REDACTED]

On 14 Feb 2013, at 3:42 PM, Joscha Bach wrote:

> Dear Stephen,
>
> thank you for your so far unending patience in that discussion. -->
> =tp://www.xkcd.com/386/

That's a terrific website!

>
>> 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 =enabler 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
> =een neglected by what we now often call "classic AI". (But if we look
> =t the original ideas, that was probably not intentional. Logic based
> =ystems were low hanging fruit. Once you have a paradigm and a
> =community, you end up with a methodology that is bound to stay,
> =nfortunately.)

That might be true.. but I think LISP lent itself to a certain way of thinking....

- >
- > 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 think that the most interesting difference between chimps and humans is how we can use grammatical language to "program" and organize our thought processes, and how we can use it to suggest, transmit, create and manipulate new categories.

Perhaps. I'm not convinced.

>

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

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.

>

>>>

>>> --]]]]]]]]]]]]]]]] So.. what would be wrong with building a machine that could do well on IQ tests?

>>

>> Nothing is wrong with building a machine that excels at playing chess or cooking coffee or scoring the Raven test.

>>

>> The Raven would be a bad idea -- way too easy. 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

> - The working memory tests are likewise rather simple engineering

> =roblems

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

> - 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 as 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 =omputer to do well on IQ tests. I would love to see this!

>

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

Yes

>

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

I think most of the above is in fact implicit in some of the tests. =emember that factor analysis reveals a very rich structure of human =ntelligence, with 60+ specific identifiable abilities that feed into =t.

>

> 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 =raining, used me as a guinea pig for testing -- and I found taking the =est really interesting... and was surprised by what I found trivially =easy and what I found more challenging)

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

>

> While the literal understanding of the Turing Test leads nowhere (or,

> =ell, to the Loebner prize), I think that he had the right idea.

> =ntelligence is reflected in the ability to participate in meaningful

> =iscourse,

It may be reflected, but such discourse is not a necessary consequence =f intelligence. A deaf mute could still be very intelligent.

> which includes interpreting and creatively structuring the world. Many =f the things that the WAIS measures, like recognizing and categorizing =hapes, are prerequisites for that. Others might be acquired tastes that =merge on more basic functionality, like mental arithmetic. But a =oolbox 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, =hich could in principle arise from numerous architectures..

>

> Some of the basics stem directly from the requirements of producing =dequate representations of perceptual and abstracted content (hybrid =ierarchical representations that can do associations, compositional =tuff, grammatical systematicity, learning and categorization, =nheritance etc.). Others come from the needs to get the processes of =ottom-up/top-down perception, reflection, memory retrieval, inference, =nalogy building etc. to work. And some have to do with the requirements =f translating between Mentalese (in Pinker's sense, not in Fodor's) and =atural language.

I don't disagree with any of the above

>

> I may delude myself in thinking that I know what the basics are. In =act, it is extremely likely that I do (every computer science problem =eems to be misconceptualized until it has been properly implemented).

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

> 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.

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

>
>> 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.

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

>
>> 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 Marvin 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.

I believe that Richard Gregory had a version of that argument well before Sloman

> Damasio's argument does not convince me because he does not elucidate a functional role that would make emotion an absolute requirement for an artificial mind.

I think emotion serves to prioritize goals, which seems kind of important

> 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.

>
Lesion data were the beginning, but there have been many sorts of studies now of the role of emotion in reasoning. The "social neuroscience" folks have done a lot of this stuff. I'm convinced.

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

>
I think so... I recall Kevin Ochsner years ago telling me how lousy the card test was, and that there were much better ways to get at the same point (but I no longer recall what those studies were..)

> Please do not misunderstand me; I am making a merely philosophical point here, with respect to the basic requirements for Intelligence. I think that emotion is highly interesting, that Damasio is quite correct with respect to

what emotion does, and that it makes a lot of sense (and =s fun) to equip AIs with emotion, mood, affect and emotional =ispositions. But strictly necessary? No.

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

>

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

>

> 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)

I don't know about this

> , BICA (a remnant from the failed DARPA proposal of the same name, and

> =ater picked up by Alexei Samsonovich as an alternative to AGI, I

> =uspect because he does not get along with Ben)

Yes, the history is correct -- I was part of one of the original teams.

> , and Cognitive Systems (Pat Langley et al.). I basically like them

> =ll, and think that they should join forces, while simultaneously

> =aising the bars against narrow AI and science fiction. Many members

> of =he audience already belong to two or even three of the groups.

> Alas, =olitics, mutual accusations of scruffiness and stuffiness, and

> so on...

Weird. I had no idea that AGI or "Cognitive Systems" existed... Another =easure of how out of touch I've become...

>

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

>

They have a journal now; it might be worth a glance..

Be well!

s.

> 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>1360902239</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>276496</string>
</dict>
</plist>
```