From: S.M. Kosslyn <[redacted]
Sent: Thursday, February 14, 2013 1:33 AM
To: Joscha Bach
Cc: Jeffrey Epstein
Subject: Re: Today's discussion

Hi J..

I do indeed use the Mail app on a Mac. So.. I'll just write, and see if =t does the job!

Thanks,

s.
_/

Stephen M. Kosslyn
Director, Center for Advanced Study in the Behavioral Sciences Stanford University
75 Alta Road
Stanford, CA 94305
Voice: 1[redacted]

On 13 Feb 2013, at 12:13 PM, Joscha Bach wrote:

> Hi Stephen,
>
>> OK, I'll risk yet another level of embedding... please see
>> =-]]]]]]]]]]]]]]]]
>
> my computer thinks that you are using Apple Mail. It can take care of =he embedding automatically, by adding another layer of embeddementation =or each response iteration. Usually, you can simply write your answer =ithout any embedding, so newer parts of the conversation can be =ecognized by having fewer indentations.
>
>>> ... could also be applied to a lot of other, much dumber or =on-social animals.
>>
>> --]]]]]]]]]]]]]]]] Yes, and they have less of it (but nevertheless
>> =ts the same "it")
>
> I agree, and yet only a subset of animals can use symbols for =ommunication.

communication and intelligence are not the same thing (think about bees =tc)

> 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.

>
>> --]]]]]]]]]]]]]]] 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)

> But these skills are not sufficient preconditions for Intelligence.

I'm not convinced. I think the skills necessary to do the WAIS are in =act those on which intelligence is based -- and there's quite a bit of =vidence to support this intuition.

> If we want to build a bootstrapping mind (i.e. a toddler level =ntelligence with the capability to learn language, form and use =dvanced abstractions etc.) then starting with chess or Raven is likely =o result in yet another electronic idiot savant.
>
Not so clear to me

> I think that minds are architectures with representations + cognitive =ools.

Absolutely! And that's just what is needed to do well on a WAIS

> They are organized and equipped to learn how to be get some =roficiency at chess or other puzzle tasks, like IQ test suites. Once we =ail the basics, we can probably scale them way beyond human capability, =ecause hardware does not suffer from the same resource constraints as =etware.

How do you know for sure what the "basics" are? I think one approach is =o analyze what abilities and skills are needed to accomplish a range of =asks -- and the WAIS presents such a range of tasks, which are designed =o tap different abilities and skills...
>
> Going for the puzzles first is likely to bypass the subtleties of =niversal mental representations, language etc., that are prerequisites =f Intelligence.
>
I don't follow this reasoning

> More specifically, I would try to address a test like the Raven

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 =ypes of items -- solved by spatial vs. analytic strategies). The Raven =oes not even begin to characterize all of what is captured by the WAIS

> with training a bunch of nested classifiers for the visual input, and =onnecting these to the equivalent of Hofstadter's Slipnet (see "Copycat =rchitecture"). We could probably bruteforce a distance metric for the =ransformations then. It is likely that I will hit a few bumps on the =oad, but once it works, we would have another narrow AI classifier =odel. Tests that mix verbal and arithmetic performance with geometric =equence tasks will require a bigger library of tools. Perhaps we will =ven need to scale it up to a Watson level project. And yet, it won't be =trong AI.

I agree -- for the Raven

>
>>>> ====>>There are many reasons why ACT-R falls short.

>> --]]]]]]]]]]]]]]] I think the SOAR community is even worse
>
> I don't think so!
>
> First of all, the Soar guys never said that their architecture exactly
> =atches what the brain does, down to neural firing rates and the
> wiring =f the basal ganglia. Being a Soar acolyte does not require you
> to =ubscribe to the creed that there is some kind of 1:1 mapping
> between =oar and neurophysiology, which means that you are free to
> add, alter =nd remove functionality without making implicit statements
> about brain =natomy. (In Act-R, only John Anderson and his inner
> circle can make =cceptable changes to the architecture itself.)
>
> Next, there is no tightly woven community. John Laird has given up on =ard AI (although he wants to do it, he does not see a good way to =roceed), and Paul Rosenbloom has started a new architecture, after ten =ears of abstinence.
>
> The only bad thing that I would say about Soar is that it really does
> =ot care all that much about being a good model of the mind,

Right

> is mostly an enhanced edition of "General Problem Solver Strikes Back"

Right

> , and with its focus on applications, has partially turned into a =lassical narrow AI paradigm.

Right.. and it constrains its users in too many ways
>
>> --]]]]]]]]]]]]]]] A.Einstein's theory really got traction when it
>> =redicted phenomena (e.g., light bending around gravity wells) that
>> no =ther theory predicted
>
> Absolutely. But while answering the big questions, he himself did not =eem to care so much; to my knowledge, he has never conducted a single =xperiment (apart from trying to see with how much private hobby time =he patent office would let him get away with).

He was a theoretician, not an experimentalist -- but he made it very =lear what were the novel predictions of his theories (and left it to =thers to test them, which is fine).

>
>> --]]]]]]]]]]]]]]] I agree completely, if you include emotional
>> =eactions as part of the thought process
>
> I agree. But I am not convinced that proper emotions are absolutely =ecessary for Intelligence (motivation might suffice to drive some kind =f non-emotional, serene Buddha intelligence). I am nevertheless =nterested in understanding and modeling them.

I think Antonio Damasio and his successors have made a very good case =hat emotion plays a key role in reasoning. (Not just motivation, actual =motion.)
>
>>> Again, thank you for the time and care that you take for responding!

>>> =>>
>> --]]]]]]]]]]]]]]] My pleasure! I'm finding these interactions very =timulating.
>
> Same here! It is really nice to meet someone actually interested (and =f course extremely knowledgeable) in these topics!

Thanks.. But I am a bit "out of it"... it's been several years since ='ve been to a relevant meeting or even had much interaction on these =opics (and I stopped reading the journals at least six years ago). Are =ou involved in BICA? That seems like a natural community for you!

Be well,

s.

> Bests
>
> 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>1360805565</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>276270</string>
</dict>
</plist>
```