
From: Joscha Bach [REDACTED] >
Sent: Wednesday, October 23, 2013 3:10 PM
To: Sebastian Seung
Cc: Joi Ito; takashi ikegami; Ari Gesher; Kevin Slavin; Martin Nowak; Greg Borenstein; Jeffrey Epstein
Subject: Re: MDF

Hi, Sebastian, what an honor to meet you!

> I call this theory the "hierarchical perceptron," and credit it to
> Fukushima's Neocognitron (1980). His work in turn was inspired by
> discoveries and speculations of neuroscientists Hubel and Wiesel
> (1962). Today's deep learning architectures for vision continue in
> =his tradition.

By the way: Kurzweil's most recent book mostly covers his rediscovery =f the Neocognitron, too.
I believe that it is a crucial part of the story of the mind (with =espect to hierarchical self-organization), but not the whole story, of =ourse.

> In the 1990s, my research was focused on machine learning, which can
> =e seen as a rebranding of the pattern recognition camp of AI. As far
> as = can tell, AGI is an attempt to revive and rebrand the reasoning
> camp =f AI. I'm sympathetic to this goal. If I were starting over in
> AI today, I'd study reasoning rather than join the pattern recognition
> bandwagon.

I perceive AGI (including BICA, cognitive systems etc.) as a quite =eterogenous bunch of approaches, which are only united by their =ommitment to the original goal of AI, i.e. building/understanding minds =nstead of building applications. Most of the people in the field would =cknowledge the need to combine associative, parallel, distributed =processing with some localist processes that facilitate language, =nalytical deduction, planning and so on. We need to to all of it, in a =common framework.

> That being said, AGI will have trouble succeeding because it is
> following the scruffy tradition. Perhaps the main failing of this
> tradition is its refusal to define objective (and preferably
> quantitative) measures of success.

The question of good benchmark tasks is haunting AI since its inception. =sually, when we identify a task that requires intelligence in humans =playing chess or soccer or Jeopardy, driving a car etc.) we end up with = kind of very smart (chess-playing, car-driving) toaster. That being =aid, AI was always very fruitful in the sense that it arguably was the =ost productive and useful field of computer science, even if it fell =hort of its lofty goals. Without a commitment to understanding =ntelligence and mind itself, AI as a discipline may be doomed, because =t will lose the cohesion and direction of a common goal.

> Minsky is bitter because he failed to turn his field into a science.

The field is there alright, still feeding truckloads of tenured =rofessors, huge conferences and so on. I think that Minsky's =frustration is much more fueled by the disregard of the contemporaries =or the "right" ideas, and the deterioration into

methodologism, =intellectual cowardice and the focus on short term goals (in some sense: =recisely what it means to become a science).

> I'm reminded of the saying that "A science is any discipline in which
> the fool of this generation can go beyond the =oint reached by the
> genius of the last generation."

You can stand on the shoulders of giants, or on a very big pile of =warfs: it works either way. But we need to pick the right questions. (I =m currently quite enamored with a developmental approach, i.e. working =long child performance with respect to story and scene comprehension.)

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>270440</integer>
    <key>date-last-viewed</key>
    <integer>0</integer>
    <key>date-received</key>
    <integer>1382541011</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>354477</string>
</dict>
</plist>
```