

From: Joscha Bach <[REDACTED]>
To: Jeffrey Epstein <jeevacation@gmail.com>
Subject: Re:
Date: Fri, 09 Mar 2018 05:05:59 +0000

Last week I got to know Steve Hyman, Daniel Kahneman and Bob Horvitz. Telefonica invited all of us to a two day workshop with Pablo Rodriguez, Ken Morse and a few others, where we were meant to advise them on how to use AI for health applications. I told them that I think the goal of therapeutic invention is not to increase happiness, but integrity. Happiness is merely an indicator, not the benchmark. Current apps tend to subvert the motivation of people, but I don't think that this is necessary or the best strategy. Humans are meant to be programmable, not subverted. They perceive their programming as "higher purpose". If we can come from the top, supporting purpose, instead of from the bottom, subverting attention, we might be more successful.

(Downside might be that we create cults.)

Of the bunch, Hyman managed to be the most interesting (Kahneman was very charismatic but mostly tried to see if he could identify an application for his system one/system two theory). Gary Marcus was there, too, but annoyed everyone by being too insecure to deal with his incompetence.

Did I tell you that I discovered that Deep Learning might be best understood as Second order AI?

First order AI was the classical AI that was started by Marvin Minsky in the 1950ies, and it worked by figuring out how we (or an abstract system) can perform a task that requires intelligence, and then implementing that algorithm directly. It yielded most of the progress we saw until recently: chess programs, data bases, language parsers etc.

Second order AI does not implement the functionality directly, but we write the algorithms that figure out the functionality by themselves. Second order AI is automated function approximation. Learning has existed for a long time in AI of course, but Deep Learning means compositional function approximation.

Our current approximator paradigm is mostly the neural network, i.e. chained normalized weighted sums of real values that we adapt by changing the weights with stochastic gradient descent, using the chain rule. This works well for linear algebra and the fat end of compact polynomials, but it does not work well for conditional loops, recursion and many other constructs that we might want to learn. Ultimately, we want to learn any kind of algorithm that runs efficiently on the available hardware.

Neural network learning is very slow. The different learning algorithms are quite similar in the amount of structure they can squeeze out of the same training data, but they need far more passes over the data than our nervous system.

The solution might be meta learning: we write algorithms that learn how to create learning algorithms. Evolution is meta learning. Meta learning is going to be third order AI and perhaps trigger a similar wave as deep learning.

I intend to visit NYC for a workshop at NYU on the weekend of the 16th.

We just moved into a new apartment; the previous one had only two bedrooms and this one has three, so I can have a study. It seems that we are as lucky with the new landlords as with the previous ones.

Bests, and thank you for everything!

Joscha

> On Mar 8, 2018, at 16:37, jeffrey E. <jeevacation@gmail.com> wrote:
>
> progress?

>
> --
> please note
> The information contained in this communication is
> confidential, may be attorney-client privileged, may
> constitute inside information, and is intended only for
> the use of the addressee. It is the property of
> JEE
> Unauthorized use, disclosure or copying of this
> communication or any part thereof is strictly prohibited
> and may be unlawful. If you have received this
> communication in error, please notify us immediately by
> return e-mail or by e-mail to jeevacation@gmail.com, and
> destroy this communication and all copies thereof,
> including all attachments. copyright -all rights reserved