From: Ben Goertzel <[redacted]>
To: "jeffrey E." <jeevacation@gmail.com>
Subject: Re:
Date: Sun, 14 Aug 2016 18:59:30 +0000
Attachments: Prime_AGI_HighLevelPlan.pdf

Hi Jeffrey,

good to hear from you -- ope you're well !

I previously sent you a video of OpenCog doing some simple logic inference via the Hanson robot head — nothing deep inference-wise, but a software-dev milestone for us in terms of systems integration (language generation, language comprehension, speech, inference, etc. all working together OK…).

https://www.youtube.com/watch?v=LduD7Et_cOs

We are aiming for some additional videos showing more OpenCog stuff integrated w/ the Hanson robot head by early September…

When you have time we can Skype again …

Also I will be happy to send you a copy of my book "The AGI Revolution"

https://www.amazon.com/AGI-Revolution-Artificial-General-Intelligence/dp/0692756876

if you remind me the best address… Some bits of it may be too lightweight for you, but there are some interesting conceptual/theoretical sections too…

I am in Seattle visiting my mom and sister now... I'll be in San Francisco Aug 22-24 then head back toward Asia...

LANGUAGE, MUSIC, AGI ARCHITECTURE, ETC.

After our last meeting w/ you, Linas and I talked more about music perception and music learning, and how it ties in with language…

In both cases, obviously, there's a grammar part (for music the grammar is the stuff classical music theory deals with — chords and scales and harmonies and such) … and then there's a timing and gradual-change part (in language this is prosody, pauses, intonation, etc. — which is key stuff for child language learning, and helps bind language to nonlinguistic perception, etc.)

In terms of AGI architecture, on the face of it this suggests a neural-symbolic architecture of sorts would be helpful…

Symbolic methods are natural for grammars … rules of linguistic grammar or musical grammar can be learned (or programmed) as formal

structures, and manipulated by "logic rules" of various sorts…

Sub symbolic methods like the currently fashionable "deep neural nets" should be good at learning continuous-variable stuff like timing and emphasis… The emotional dynamics of timing as related to arousal and frustration, as I mentioned in an email to you earlier (the one that made you suggest I was stoned ;), would "straightforwardly" (but not trivially) emerge from deep reinforcement learning methods….

I have thought a bunch about how to embed deep neural nets into an architecture like OpenCog, so as to enable feedback between symbolic and sub symbolic operations… Music as well as the expressive/continuous-variable aspects of language would be fascinating in that regard…

GOOGLE…

I spent a day visiting friends at Google earlier this week…

Smart people and some interesting projects, but I feel fairly confident there is no actual "AGI system building" going on at that place right now….

Kurzweil's team is going nowhere — all the researchers on his time quit and he's just working on a souped-up chatbot.

DeepMind is far more serious but they're doing a mix of academic paper publishing, awesome demos and helping various Google divisions with machine learning …

Understanding the cost structure of doing stuff within Google was also instructive for me… A team of 5 guys within Google costs them about $2M/ per year all considered…. Whoa….

THE REQUISITE BORING PRACTICAL STUF…

At the moment, as, you know, our main source of funding for OpenCog is Hanson Robotics (this is what pays my salary for example), but this is very tenuous as they are a startup constantly on the verge of running out of money ;p …. Jim Rutt is funding Nil Geisweiller's work on PLN logical inference which is great…

The plan Jim and Cassio and I worked on, which I sent you before (and reattach here for your amusement), would cost $2M per year for 3 years if fully funded — which is equivalent in cost to a typical small-team project at Google (probably 1/5 what AlphaGo cost to DeepMind, and 1/2 what the Atari 2600 demo cost DeepMind before that)….

Basically this plan would allow what we discussed when Ehud Barak was at your place. An AI that we could teach arithmetic, and so many other things, to via directly interacting with it in the physical world. This would enable it to learn so many things in a grounded way….

Obviously any help you could provide toward this would be awesome. But I understand you've got loads of financial demands and also lots of other interesting projects on your horizon… so it goes…. I will keep progressing as best I can... and when I finally get to a really amazing demo I will obviously let you know...

Guess that's enough for now....

-- Ben

On Thu, Aug 11, 2016 at 8:47 AM, jeffrey E. <jeevacation@gmail.com> wrote:
> News?
>
> --
> 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
>

--
Ben Goertzel, PhD
http://goertzel.org

Super-benevolent super-intelligence is the thought the Global Brain is currently struggling to form...