Hi Jeffrey,

Good to hear from you -- and timely, in fact! I was planning to email you in a few weeks time, but since you've pinged me now, I'll take the opportunity to tell you now all the stuff I was going to tell you then....

There is a LOT to say, so please read to the end ;-) ...

In case the email formatting is annoying, I've attached the same message in PDF form as well...

I'll start with updates, and then move on to hopeful near-future stuff...

On a personal note: I'm living in Hong Kong these days (trying to short-sell my empty house in Maryland, whose mortgage is currently badly underwater), and have moved on to my 3rd wife (a very sweet 28 year old mainland-Chinese computational linguist, about to finish her PhD).

The "OpenCog for game characters" project (based at Hong Kong Polytechnic University) is proceeding OK. It was supposed to end January 2013, but staffing the project took longer than expected, so there's extra $$ in the project account and we will probably be able to get it extended till June. As you may recall, that project was funded 9/10 by the Hong Kong government (specifically the ITF -- "Innovation in Technology Foundation"), and 1/10 by my consulting company Novamente LLC (which got the money indirectly from you, via your donation to the nonprofit org Humanity+).

I have been working on the OpenCog AGI project part-time (as there's not enough $$ in that project's budget to pay me a salary), and there have been 4-7 people working on it full-time (some are MS or PhD students, a couple are young programmers). I'm also involved with a couple other AI-based projects here in Hong Kong that pay me money (not a lot, but enough to let me pay my kids' college tuition in the US, etc.). These other projects involve using OpenCog's machine learning software to analyze genetics data, and Hong Kong stock market data. But I go to the OpenCog AGI lab for a few hours every day, and spend a lot of time on it evenings and weekends.

We have not yet gotten the OpenCog game character to be extremely intelligent. It's definitely still a work in progress. But we have gotten OpenCog to control a character that moves around in a game world and does various things, trying to achieve its simple goals (finding batteries, avoiding danger, learning new things, etc.). The game world we've built, in the Unity3D game engine, is blocks-based and involves 100,000+ blocks used to build houses, trees, etc. etc. Actually, just one week ago we finally got OpenCog to work together with this fancy game world -- it took too long, but it was great to see it finally happen. So our job for the next 6 months is to make it more and more intelligent.

My collaboration with Itamar somewhat stalled for a while, due to him back-burnering his AGI work in favor of commercial work (AI-based stock market

prediction). However, I now have someone here in Hong Kong (funded directly by me, outside the ITF project) working for a few months on modifying Itamar's computer vision system (DeSTIN) to make it more compatible with OpenCog, so we can use it in 2013-2014 as OpenCog's visual and auditory cortex.

Remember in 2001 you funded me to write a book on my design for AGI, and I didn't finish it? A first draft of that book is finally done ("Building Better Minds: The CogPrime Architecture for Artificial General Intelligence"); it's about 1000 pages long and you can see it online at

http://goertzel.org/monkeyburger/bbm_main.pdf

It is complete regarding content, but still needs editing for style, fixing references, aligning figures, etc. At the moment a few colleagues are going over it carefully and making a long list of things needing fixing. My plan is to send you a copy of the proofs once it's finally truly ready for the publisher. As your time is scarce, I don't want to waste your time having you read a messy version when there will be a clean version available a few months later. Just for fun, I have attached the Table of Contents to this email ...

We have gotten an offer of more funding from the Hong Kong government, for a project spanning 2013-2014, which is supposed to focus on using OpenCog to control one of David Hanson's Robokind humanoid robots:

http://hansonrobokind.comDavid Hanson has been spending about 1/4 of his time here in Hong Kong, and has become a fairly close friend. I've also spent a bunch of time with Mark Tilden, the maker of the Robo Sapien (http://en.wikipedia.org/wiki/Mark_Tilden), who is based in in Hong Kong as well.

The Robokind's hands aren't that great, but we will supply it with a bunch of foam blocks so it can do blocks-world type stuff, similar to in the video game world. Its camera eyes are high-quality, so the main aspect we'll need to deal with for the Robokind project, that we aren't needing to deal with for the current game-AI project, is machine vision. For this we will use the DeSTIN framework, currently being improved my someone working for me as I noted above.

This funding for the robotics project is of a similar nature to the game-AI funding: 9/10 Hong Kong government, 1/10 me. So I was going to email you in a few weeks to ask you if you'd be willing to help me with my 1/10 of the money like you did before. That would be very greatly appreciated, obviously. What is needed is US$20K in December, and another US$20K in late 2013.

By the way, the AGI-12 conference is at Oxford University December 8-11. This looks like it will be smaller than the previous AGI conferences, due to the location being in the UK and in the winter. But it should be a very interesting group of people -- including my robotics collaborator David Hanson, who would be a really interesting guy for you to meet. If you have time, it would be awesome if you could attend:

http://agi-conference.org/2012/

Independently of you attending, if you were able to help with a small amount of sponsorship money that would be great. A sponsorship of $3000 or $5000 would definitely help the conference pay its bills. Oxford is a swanky location but doing things there is pretty expensive.

The fact that we've gotten offered more Hong Kong ITF funding is wonderful, but, it's not really enough funding to do what needs doing. So David Hanson and I have written up a research proposal for the OpenCog/Robokind/DeSTIN work, which I attach here, along with a brief paper on the vision-cognition integration aspect. Our goal (which may be hard to achieve) is to raise $350,000 for this, in addition to the ITF funding.

Obviously, if your foundation would be willing to fund this, that would be awesome. With this money PLUS the already-offered Hong Kong ITF funding, we could make a more emphatic push toward getting OpenCog to control a Hanson Robokind in interesting ways. As your profile seems to focus on funding universities, the money could perhaps be worked via a donation to Hong Kong Polytechnic University, where the OpenCog project is currently funded.

My thinking on the robotics collaboration is as follows. First, I think that having the AGI able to ingest and interact with real-world *perceptual data* is important. An intelligent system tends to adapt itself to whatever environment and problems it is given. So if an intelligent system lives and exists and learns only in a video game world full of blocks, it's not going to self-organize the internal structures and dynamics needed to deal with the real human world. The video game context is great as a prototyping domain, to help one refine one's AGI thinking and test one's code. But it's not going to get us to human-level AGI. I don't think robotics is the *only* path to human-level and transhuman AGI, but I think it's the most clearly comprehensible path. I understand how to create AGI in a system with vaguely human-like perception and actuation, because I can use the human mind as an analogue. I understand much less well how to create other sorts of AGI systems, though I have no doubt they are possible. So I think we should start with vaguely human-like systems, which are the "easiest" to approach, and then move on from there...

Secondly, I think that -- from a purely "marketing" perspective (where I mean marketing within academia and the scientific world, as well as more broadly) -- having a cute little boy humanoid robot that can walk around and talk and interact, will do wonders for getting buy-in for AGI. This will be what I have called the "AGI Sputnik" moment (a term I coined in this somewhat rambling interview, http://hplusmagazine.com/2011/03/30/seeking-the-sputnik-of-agi/). If we can pull this off, getting AGI funding will not be a problem anymore. IBM's Watson and Google's self-driving cars have already made the world more optimistic about AI, in the last year or two. A palpably intelligent little humanoid robot can be the next

step in terms of public awareness and excitement, as well as science and engineering.

...

Finally, I can't resist the urge to end on my usual note! Remember in 2001, we discussed making a "Brooklyn Project" to create AGI (though I hadn't come up with the term "AGI" yet, then). I firmly believe that, if you had decided to go ahead with that project -- and if it had been managed well -- we would have a human-level thinking machine by now. Instead, it's 11 years later, and we do not have a human-level thinking machine yet. Why not? Two main reasons, I think:

1) AGI is hard, and nobody but me has put enough sustained thought into how to design it [because it takes a particular sort of creative and pigheaded cast of mind, which I have, to want to spend soooo many years figuring out in detail how to design something huge, in the absence of funding to actually build it. The reason I have a workable detailed AGI design and nobody else does is not because I'm smarter than them -- I'm very smart, but there are certainly other very smart folks in the AI field. The reason is that I'm very smart AND a deep thinker about the mind AND, perhaps most critically, I have consistently, over a long period of time, allocated a significant percentage of my workweek to figuring out a workable detailed AGI design. Whereas other researchers have focused their time on other things, from which they could more easily get short-term tangible rewards, like money or publications or tenure or grants....]

2) Current research funding avenues, namely academia and industry, are both badly suited for creating AGI

The case against academia as a place to get AGI built was recently articulated by my friend (the AGI researcher) Nick Cassimatis in the following interview

http://www.forbes.com/sites/markchangizi/2012/11/09/for-siris-new-competitor-skyphrase-academia-isnt-big-enough-for-ai/

A key quote:

***

Since our goal is to actually identify mechanisms that are powerful enough to achieve human-level intelligence, the best way we have of proving that our theory is correct is to actually implement it and show it actually understands language at a human level. It's actually surprisingly difficult to get research like this published and supported within normal academic communities because they are more interested with smaller, incremental results that can be precisely quantified. It is very difficult to get academic papers about complex systems published in the quantities you need to thrive in academia.

... one of the problems with academia today is that one's career progress is disproportionately linked to bringing in money (almost always government money). When one asks oneself how to best ensure getting a grant, the answer is invariably, "Keep doing more of whatever got money before."

***

Nick, in the interview, is explaining why he decided to pursue his AGI research via a new company rather than via academia (he's a tenured prof at RPI). However, he is new to the business world. He hasn't yet realized that the business world is focused on making money or getting customers, not on scientific innovation. He complains that academia is focused on "smaller, incremental results that can be precisely quantified." -- but he hasn't yet realized that business also focused on smaller, incremental results that can be precisely quantified. It's just that the incremental quantitative results business focuses on are money or customer-acquisition oriented, whereas the ones academia focuses on are stuff like "getting accuracy a few points higher on some accepted benchmark problem."

If neither academia nor industry is well set up for the creation of AGI, then how can we get it done?

One option is that some wealthy individual comes up with a design for AGI himself, and pays to get it built. Jeff Hawkins tried this, but unfortunately his idea isn't very good (http://numenta.com) ..... It's basically like DeSTIN, and will never be more than a perceptual cortex.

Another option is open-source development, like the way Linux was built. We're trying this with OpenCog (http://opencog.org). But building AGI requires people with much more specialized expertise than building a "mere" computer operating system like Linux. So getting OpenCog off the ground is going slowly. In early 2013 we will publicly release our code using OpenCog to control video game characters, and try to get open source programmers around the world psyched to help with that. But frankly, I don't think that's going to be how the core of AGI gets built. An active OpenCog open source community may be a big, big help to getting OpenCog AGI built, but my best guess is that the really critical parts of the AGI are going to be built by a small team working closely with me in a single location.

Another option is that someone like Jeffrey Epstein, decides to fund an AGi project like OpenCog.

...

So, summing up....

What would be a VERY GOOD avenue for getting from here to human-level AGI?

-- Jeffrey Epstein funds the $350K OpenCog/Robokind proposal

-- In early 2014, once there's some demonstrated success from using OpenCog to control the robot, we finally pull together multimillion dollar funding for AGI and start moving FAST toward the end goal (of human-level AGI, as a first step toward transhuman AGI)

Of course an EVEN BETTER avenue would be having multimillion dollar funding for OpenCog now. But I'm trying to balance what I think would be best; versus what I can expect you to find reasonable from your own point of view...

But if the $350K is not palatable to you, obviously I will still be very happy to have your support on the next ITF project ($20K now and $20K in late 2013), if you are kind and visionary enough to be willing to do that...

And, that about wraps it up ;)

Ah... one more thing. I will be in the US from roughly December 18 through January 15. It would be great to get together and discuss in person, if you'll be available anytime during that interval. Please let me know!

Thanks

Ben