

From: "[REDACTED] on behalf of Ben Goertzel <[REDACTED]>"

To: Jeffrey Epstein <jeevacation@gmail.com>

Subject: Reviews of CogBot proposal, at long last... ;-)

Date: Mon, 18 Jan 2010 05:57:16 +0000

Attachments: BriefResponses.pdf; CogBotReviews.pdf

Hi Jeffrey,

At long last, and with apologies for the delay, here are the reviews that you asked for.

The reviewers are:

-- Leslie Pack Kaelbling, the director of MIT's AI lab (and an AAAI Fellow) and an expert in both AI and robotics

-- Paul Rosenbloom, a senior AI prof at USC (long-time Deputy Director of USC's Information Sciences Institute, and co-developer of SOAR, perhaps the world's leading AI architecture)

-- Christian Lebiere, AI/cog-sci prof at Carnegie Mellon, and the current leader of the ACT-R cognitive architecture (perhaps second only to SOAR in reputation) as well as an innovator in neural net AI

-- Stan Franklin, leader of U. Memphis's Institute for Intelligent Systems and inventor of the LIDA AI architecture

...

Stan's comments are brief, Paul's and Christian's are fairly extensive, and Leslie's somewhere in the middle.

For your curiosity I have also included a file giving some email responses I made to Leslie and Christian in response to some minor concerns they raised in their reviews. My responses to Stan were on the phone, and my discussions with Paul were mostly F2F, so I don't have a record of my conversations with them (though in fact Paul and I had the longest and most interesting conversations on the proposal).

Obviously we hope these letters will satisfy your requirements!

None of the reviewers were able to really understand the depths and details of our AGI approach based on the proposal, which is as expected. What they were able to do is validate that the approach seems to generally make sense. For them to understand the depths and details they would need to read the book I am now writing (with help from Cassio, Nil Geisweiller and some others), and then think about it a lot....

In other words, the reviewers don't know enough to understand why I think my approach can lead to a human-level AGI. But they do know enough from the proposal to understand why it might be sensible for me to think my approach can lead to a human-level AGI ... and to think that, even if I'm overoptimistic, a lot of interesting stuff will be learned during my attempt...

FYI, these are the only letters I got -- I didn't reject any letters due to being mediocre or negative. And no one read the proposal but told me they couldn't write a letter because they didn't like it. I sent the proposal to 8 people and these are the 4 that responded with letters (Paul after I paid him a visit F2F, because I had another reason to go to LA). Of the other 4, 2 said they were too busy, 1 never replied, and 1 promised the letter but didn't send it yet....

Paul and Christian I knew slightly beforehand (due to chatting with them at conference) but they didn't know much about my work beforehand. Leslie I had never encountered before except for a brief email exchange last year in which I invited her to a conference but she declined due to being too busy. Stan I know better because we co-organized AGI-08 together,

but, he had never read so much about my research before.

I'm in China now and return home the night of Jan 28. If the letters leave you positively inclined, perhaps we could get together sometime in early February to discuss the project.

If you don't end up funding the project, it's quite possible I'll relocate to China and try to get my AGI system built here, where I have Chinese NSF funding to pay grad students. However, this is by far an inferior possibility because the work will go 3-5x faster if I can use my existing team, rather than training a bunch of Chinese in all the complex stuff my existing team knows already. If I could import half my existing team here to China with me, then I could probably get the Chinese up to speed on my project quickly and make good use of them (working alongside my existing team). But what I'm seeing on this visit to China is that, even though the Chinese grad students and young profs are really smart, my existing team has so much built-up knowledge from working on my AI system over the years, that trying to do the project without them would really be vastly inferior....

If you do choose to fund the project, one possibility would be to fund it via Itamar's university. He says he can work it out so that there is zero overhead, which is appealing. On the other hand, there's also an argument to do it via some university in the DC area where I live, which would have the property that it's easy to bring government types by the lab for demos (which is good for getting ongoing funding from government agencies). Of course, the cheapest option would be to move my team to China where I'm now sitting and fund the project here

[because salaries can be lower here due to the low cost of living]; but my feeling is that for you the precise amount of \$\$ is less a concern than your faith (or otherwise) that the project can succeed...

Anyway I'm now very curious for your reaction.... Actually the letters are about as good as one could expect to get from US academics, since they are a critical lot. At least, they are a far cry from Schank's ill-informed reaction ;-p I almost feel, upon reading the letters, that I did too good a job of making my approach seem mainstream (the proposal was quite honest, but it seems not to have highlighted my points of disagreement with the reviewers -- which is because these pertain mainly to the method of integration of different AI components and the way this integration connects to the emergent structures intended to emerge in the system's mind as it learned, which in my view is absolutely critical, but is hard to write about in a short space...)

BTW my cellphone seems to work here in China BTW, though not all calls get through, so if you wish you can call at 240 505 6518. Due to time zone differences your morning or evening is the best time to call; otherwise you may find me at a suboptimal state of lucidity....

As a side comment, I mentioned your name to Paul and Stan, but not to Christian and Leslie -- not for any strong reason ... just because I know Stan well and I met Paul F2F, whereas my interactions with Christian and Leslie were more formal and via email.

Next move is yours ;-)
ben

--
Ben Goertzel, PhD
CEO, Novamente LLC and Biomind LLC
Director of Research, Singularity Institute for AI
External Research Professor, Xiamen University, China
[REDACTED]

"I have a deep nostalgia for the future." -- Max More