

Panel Discussion: What are the bottlenecks, and how soon to AGI?

Participants: Dr. Stan FRANKLIN, Dr. Hugo DE GARIS, Dr. Sam S. ADAMS, Dr. Eric BAUM, Dr. Pei WANG, Steve GRAND, Dr. Ben GOERTZEL (Moderator: Dr. Phil GOETZ)

A video version of this dialogue is available at www.agiri.org/workshop

[Ben Goertzel]: In general, what was intended when forming the question for this discussion was - what is the feasibility of achieving AGI within a reasonable time-frame? I know we have some widely differing opinions on this. I guess we can go around and have everyone give a brief statement of their views on the topic. Our moderator will be Dr. Phil Goetz, and I would like to encourage people to allow the moderator to moderate as he is a very moderate guy.

[Phil Goetz]: Everyone will give an opening statement and then we will follow with audience questions. Ben would you like to start?

[Ben Goertzel]: In terms of what the bottlenecks are, there are two sorts of answers, and the first answer is: Both generically, and for the Novamente project in particular, the biggest bottleneck is just funding into the field. It's a common theme for many people here. For Novamente, as for most of the other AGI projects here, the main obstacle to more rapid progress is financial resources to pay people to do the work. That's very much the case for Novamente. If we had a fairly modest amount of funding we could accelerate our progress by an order of magnitude. Next, in terms of the technical bottlenecks, that gets into deeper questions. I think the nature of current programming languages is something of a pain as it slows things down a lot. We are using C++ because it's scalable and it lets us manage huge amounts of memory effectively. But we could move toward AGI a lot faster if there were a nicer programming language with anywhere near the same scalability as C++. Moving on: This is not quite a bottleneck, but I would say that if the Novamente system is going to fail to achieve AGI, which I think is quite unlikely, then it would be because of a failure in the aspect of the design wherein the different parts of the system all interact with each other dynamically, to stop each other from coming to horrible combinatorial explosions. A difficult thing is that AI is all about emergence and synergy, so that in order to really test your system, you have to test all the parts, put them together in combination, and look at the emergence effects. And that's actually hard. The most basic bottleneck is that you are building an emergent system that has to be understood and tested as a whole, rather than a system that can be implemented and tested piece by piece.

So to sum up: Once the funding bottleneck is solved, then you run against the difficulty of the problem, which I believe we have solved with Novamente design. But AGI is an empirical science and engineering discipline, so I can't yet say that for sure.. You never

know until you've implemented and tested your system in its complete form. In terms of the amount of time until AGI is completed, of course that's conditional on boring things like the level of funding and whether everyone on the team is hit by a truck and so forth. I would say, if we receive a reasonably adequate level of funding, the amount of time to thoroughly validate or refute the hypothesis that Novamente will lead to a human-level AGI, if I want to give a very wide bound, is not going to be more than 20 to 25 years, and no less than three years. I would say it's highly viable to do it within 10 years, possible to do it within five years. I would also say estimating these things is almost impossible so none of these numbers should be taken very seriously. We can't even estimate how long it will take to make a word processor or a new version of Windows. The most important thing is not time estimates, the most important thing is whether you have a design for software that actually is capable of achieving very high levels of intelligence when properly implemented and tested. I believe that we have that, and I don't think we have the only possible such design. I think there are many possible ways to do it.

[Pei Wang]: First I'll talk about my own research product. The general bottleneck I think is the time I can spend on the project. So clearly I agree with Ben Goertzel with his plans. If we can get support and more resources then we can progress faster. In my case the technical bottleneck I believe will be the inference control part. It tells about if you have limited resources in a system, how to distribute it among many activities. So that is my situation.

In general for the AI field as a whole, the main bottleneck is funding of course. Also I would like to see more people get involved in this research. We are still a very small minority in AI or Cognitive Science. Most people are working on narrowly defined stuff. A few years ago I think this kind of discussion would have been taken by many people as a joke. In the last two years we have seen a lot of change in attitude. As a whole we should try to build our community and encourage more people to do more discussion and publications. Even though we have different opinions on how to do AI, I hope we recognize something in common in this room.

Another bottleneck of the whole field of research is that we all have very different opinion about what AGI means. So ultimately I can only talk about myself.

NARS should be finished, given my current progress, within two or three years. After that, optional extensions will depend on the resources question again or if other people want to get involved.

[Eric Baum]: So I think the hard thing in building an AGI is the generalization. If you build a system, unless it is extremely special, it's not going to generalize. Just adding more and more stuff on top of it doesn't help. I don't think. I think that the way we got our intelligence is that evolution built a lot of specific code that understands and exploits real structure in the world, and finding that code is very hard. So I am skeptical about most other approaches that I've heard here. I don't think they are going in the right direction. It's hard for me to understand how you could think that, if you put one inference engine and another engine together and you load in WordNet, that's going to understand anything. I think that in humans, pretty much every word corresponds to a Python module that knows how to do structure and knows how to do stuff. I am trying to construct some of these for just Sokoban, that understand in a limited domain how they deal with topology. It's pretty hard. So I don't think most of the efforts are going in the right direction from my perspective. I think if we had a Cyc like program with millions of dollars of funding for ten

years, five years may be, we may produce enough stuff so that it looked pretty interesting; but the other thing is, I don't think that's intelligence. I don't think you are going to hit a Singularity and its going to take off, I think it builds and builds and will be getting better slowly for a very, long time. So I'm not worried about hitting a Singularity where it suddenly takes over the world.

[Sam Adams]: Timescale and obstacles. I will talk about obstacles first. Hardware is not an obstacle. We have more hardware now than we know what to do with. I will let you in on a little secrete. In the year 2000, Paul Horn, Senior VP of IBM Research, commissioned a study. We went off and looked at all the Moore's Law like things and the result was we don't have to worry, as developers. The problem, as a hardware company, we said, was what the heck we are going to do with all that stuff? What are people going to buy it for? There are two answers. High fidelity physics simulations, which can never get enough fidelity, and real AI. We got the stuff coming. So it's how we use it that is the big barrier and that's what we are after. Fidelity, we talked about embodied systems. Embodied system situated in environments. The fidelity of our bodies, hardware or software, is paltry and nowhere near what I think are necessary or sufficient conditions to get human-like intelligence. Just because we are dealing with systems that have at best tens of sensors instead of a few billion. We don't need a billion but we could go a lot richer. I mean the simulator system I'm working on, and AGI-SIM, they're all fairly simplistic and they need to be a whole lot richer. Training time is another big one. You start talking about developmental systems that are going to do their own learning, then you've got to figure out how the heck you are going to train them fast enough. Now I presented yesterday that we believe we have a way to do that with some engineering tricks and taking advantage of social tricks like the internet and multi-user games, but that's a real problem. If you are going to do developmental robotics or developmental intelligence, you've got to be able to say how are you going to train it fast, and that's a limitation. Shared resources and minds are a barrier. Probably in no other field, and those of you in other fields may strenuously argue with me, is there more "NIH" (not invented here) than in the computer science field. Frankly if we are all determined to keep stepping on each others toes, we won't get there until one brain can figure it all out. We all have separate projects -- even though they have similar architectures, they are incompatible, we could not fuse them if we tried. The same problem hit AI in the early days. Everyone had a system; no one could do anything with it. That's a serious problem. If we as a community want to solve the problem, that's the limitation we need to address. Funding being a limitation? I don't think so. I see this as someone who's job it is to convince people of the value of technology so they will invest. That's my job at IBM fundamentally. And the reason is you have to build something compelling to convince them of the value. As long as you are building toy problems, no one is going to care, no one is going to believe us, at least more than once. And guess what? The previous AI Summer and Winter already used up all the toy problems. None of them are interesting. Time frame? I believe that we are looking at serious contenders in 10 years. But they are going to have all kinds of interesting limitations, but people will definitely be sitting up and noticing and saying, wow we are really, if we are not there, we are awfully close. The common man will be able to say wow we are almost there. Once that happens, even get close to that, then the flood gates of money start opening, because everybody wants to own it. So then it will accelerate.

[Hugo de Garis]: I see three major approaches to producing AI, or as I call them *Artilects*. One could be a slave, maybe slave is the wrong word, a slave to neuroscience. Common sense says if you imitate the brain as closely as possible or sufficiently, you will end up with a conscious intelligent creature because we have ourselves as existence proof that its possible to put molecules together to make such creatures. So the brain mimic approach. Second one is just an engineering approach. Don't be a slave to neuroscience, just do whatever you like. And the third one, of course, is just the hybrid. Probably most people will do both. For example, I'm a little surprised at this workshop that there is not a larger contingent of the brain modeling crowd. I hope we will by next year's conference, or workshop. For example, Henry Markram, the Swiss IBM, the Blue Brain project, that kind of thing. So I hope at the next workshop there will be a bigger contingent of that aspect. Obstacles? Well you just talked about AI Winters and Summers. I don't know how many cycles it's been through. Is it three, five? Is there a lesson to be learnt there? I guess one obvious one is just ignorance. Even if you take a purely engineering approach, how the hell do you build an intelligent machine? My sense is we don't even know what the target is; we don't even know how difficult it is to do that thing. So it's very hard to figure out how long it's going to take to get there. Basically if you take the brain mimicking route, we will get there. Time scale? Well again, just my own opinion, I think humanity will have to have full knowledge of nanotech to enable us to create tools, sufficient tools, to be able to figure out how the brain works. Because we are talking – how many synapses do we have? What, a quadrillion? This is a huge tremendous problem. But with molecule scale tools, then I can imagine, let's say in the 2020's or that kind of time frame, then nanotechnology will explode; and then once the tools are there, and lets say in the 2030's and 2040's, there will be a corresponding explosion in our knowledge of the principles of how the brain works. As soon as some newer science principle is discovered, immediately it will go into the engineering, to the point where the two fields will just merge, they'll wed. The newer engineers and newer scientists will become effectively the same thing. So we'll have real AI's, let's say near human-level, second half of this century, with that logic. Now, if Ben Goertzel's claim is true – I mean, I'm open. I'm open to the idea. Maybe I'm being conservative and wrong on this. Maybe a purely engineering approach can correct the problem, and we haven't been able to do it till now because we just haven't had the hardware. Right? Moore's Law hasn't allowed it. But think about it: 10, 15 years into the future, we are in for an explosive increase. It's just the math. The size of the Moore doublings now is just huge. You know the argument about the lily pond. You have a bunch of lilies and I keep doubling their size every year. For many years it's just negligible. The size of the pond is huge. But the second to last doubling its sort of like half of the pond, and the last one is "boom", its there. Most of the increase will occur in the next few years. We will see it easily. I really agree, hardware is not the issue. We will have more hardware than we have ideas. My gut feeling, that's all it is, is that we are not going to have those ideas until we get them from neuroscience. But that might be conservative. And I'm open, you know, as the engineers; the engineering approach comes up with a way of doing it, because we have the hardware, well you know, that would be great. So okay: optimistically 10, 15 years -- or more realistically in my view, 50 plus; between 50 and 100, definitely this century. That's scary.

[Stan Franklin]: The disadvantage of being close to the end of an eight person panel is that essentially everything I wanted to say has already been said. Maybe not everything.

One of the things that struck me about this workshop, as other people have said -- Ben Goertzel has said it, Sam Adams has said it -- is that there is a fair amount of commonality that I see. Maybe even more than I expected, but that's not true, because we are working on the same problem, and we are going to have the same issues. So it's not surprising that solutions look much alike in ways even though they can be very different otherwise. But it strikes me that a real drawback here is the lack of a common ontology and the lack of a commonly used vocabulary. A commonly understood vocabulary, to use in discussing things with others. I found over and over again in my interactions with individuals here that our conversations could have been twice as productive in the same amount of time if we hadn't had to spend time getting around the vocabulary issue. The ontology issue. So let me say, I've had that same experience in dealing with psychologists and with other computer scientist. I don't think this is that unique to us in any way. The psychological memory researchers still have no good vocabulary for what memory is about, and they argue with themselves. Let me not even mention the philosophers. So commonality is one thing, vocabulary ontology is another, and the third is this business of funding, which everybody has talked about, and certainly I suffer from it like everybody else, and I think it's a major bottleneck. What can we do? One of the things that I would hope that we can do that I see sort of hopefully is starting here an AGI community that can interact with some regularity and try to overcome some of these bottlenecks. Then I think we should cut this thing short enough to have some discussion this afternoon about how to do that. In any event, I think that that's a major issue -- how to get this community going, because the community can do something about the vocabulary, the -- and the community may even be able to do something about funding. We need to think of ways to do better PR. As Steve Grand pointed out to me before, a lot of these bottlenecks, such as the funding issue, are because we are not presenting -- we and lot of others as well - aren't presenting a good enough case to the government, or maybe it's not even to the government. He pointed out that physicists get large amounts of money. If we had the money for the whole crew of us for one cyclotron, or one large telescope, it would -- how long would it keep us? A decade? I don't know. I haven't done the arithmetic. Okay in terms of time frame, again as Ben Goertzel said earlier, testing parts of the system doesn't work, and in an AGI agent that won't do. We have to have full systems and test them. In my view, you are not going to be able to build an AGI agent from scratch, it's just too complex and what you are going to have to do is build in whatever you can build in and let it learn. And to test for learning things takes a long time. Look how long it takes for a human child to learn to get to the point that we would call it really intelligent. I mean it is certainly some years to do that. So the whole business of testing these agents is going to be very, very long. So I am not optimistic about the length of time that it will take. Although -- yeah, I am just too scared about predictions to make any, so I think I am just not going to.

[Steve Grand]: Well, I am tempted to say that it is *all* a bottleneck. If there was a panel like this in 1960 that was asked the very opposite question: which aspects of AGI do you think are really *easy*, what answers would have been given? Those would be my answers for the really hard aspects. I'm talking about things like vision and motor control, sensory stuff and perception. Rodney Brooks has a copy of a memo from Marvin Minsky, in which he suggested charging an undergraduate for a summer project with the task of solving vision. I don't know where that undergraduate is now, but I guess he hasn't finished yet. So these are really, really hard problems. But they are the things that all of us here, or

most of us, not including me, are sort of brushing under the carpet. I think since this is primarily a symbolic AI conference, the problem we have is extending symbolic AI downwards far enough. Meanwhile, bottom-up AI people have the opposite problem - extending upwards. But you have to do it. Imagine that you were born unable to see, unable to hear, unable to feel, unable to move. How smart would you be? I think you wouldn't be smart at all, and intelligence *requires* you to have these sensory motor skills and extreme competence in perception. Yet we just don't know how to do those things. And where we think we do know, they don't fit with symbolic systems. So there's a real interface problem, and that's because I don't think we know the appropriate representation yet. So that's why I say all of it is a bottleneck; we just don't know how to do it and I think we need a whole new paradigm; a new set of representations. At all levels in the hierarchy. As for "timing" -- Alan Turing famously made a prediction 50 years ago, that by the end of the century we'd have human-level intelligence. His prediction was proven wrong at the very second that computers were about to demonstrate how stupid they are, because they couldn't add one to 1999 and get 2000. So I am not going to be drawn into this and make the same mistake. Except that I think that we really do need a breakthrough, and the thing about breakthroughs is that you can't say when they are going to come. You can't say you are halfway towards a breakthrough. It doesn't make any sense. It just happens. What we need is a paradigm shift. You can see signs that there are the beginnings of a paradigm shift: particularly in the neurosciences; particularly as new technology is coming to neuroscience that enables us to see stuff we couldn't see before. So there is a hint of a shift. But when it's going to click and somebody is going to have that "aha" moment, I don't know. It could be tomorrow or it could be in 100 years time. I really couldn't guess.

[Phil Goetz]: First question for the panel. Is AI a science? I note that we don't have a single journal that all of us reads. When we do read journals, we very rarely see the case that someone frames a hypothesis and a null hypothesis and has a statistical test to distinguish between them. We don't have repeatable experiments because no one in the field can repeat the experiments of anyone else in the field really. We don't have post docs, we don't have degree programs, we don't have a curriculum, we don't even have an AI prize, that's the equivalent of the Nobel prize or something.

[Ben Goertzel]: There's the Loebner prize.

[Phil Goetz]: But nobody's won it. Do we even really have a community in the way that many other sciences do? So I guess I would like to pose the question to whoever among you would choose to answer it: Have we reached the status of being a science and if not, what can we do to be more scientific?

[Steve Grand]: We've never really been a science. AI has always been an engineering discipline. Engineering tends to draw upon science rather than contribute to it. And as for community: yeah, one of our problems in being a community is that we all hate each other's guts. We are forever dividing ourselves up, rather than coming together, which is probably a sign that we don't have a good theory yet.

[Audience]: Ben, you seem more optimistic. Could you talk about your perspective?

[Ben Goertzel]: Well I have a quite different opinion than that of Steve Grand in that I don't think an amazing conceptual breakthrough on the level of the discovery of the quantum or curved 4D space-time, or something like that, is needed to create general intelligence. It might be needed to create provably stable friendly AI, like Eliezer Yudkowsky would like. I tend to think of the brain as a complex system composed of a

bunch of evolved kluges for solving particular problems, which have been hacked together and adapted by evolution. I think if you assemble subcomponents solving the appropriate set of specialized problems, as well as a fairly weak general problem solver, and they are hooked together in a knowledge representation that works for all the components, with learning mechanisms that let each component learn from each other -- then you are going to have an intelligent mind that can be taught. You make it a baby and you teach it, then it grows up. I don't believe an amazing breakthrough is needed, I just think we have not done what I have just said, in part because hardware hasn't been there, and part because saying it is a lot simpler than doing it, because what I just said involves a whole bunch of different components -- each of which has to basically work, with a lot of testing and tuning of the parameters of the learning algorithms. I agree that it's an engineering problem, and it's a big and hard engineering problem, but I am not sure it requires an amazing conceptual breakthrough in order to do it. You know there is a lot of stuff in the Novamente design that I think are moderate level conceptual or scientific breakthroughs, which are really interesting stuff, as a single example, the stuff Moshe Looks is doing. I think Moshe's work shows that we have figured out how to make evolutionary learning work, which is one among many mid-sized breakthroughs we've had to make as part of creating the integrative Novamente system. But none of those things is an amazing epochal breakthrough, and I don't think one is needed.

[Steve Grand]: I wasn't actually suggesting that we needed a magical breakthrough, just a breakthrough, and it may be very simple: a change in the way we look at the problem.

[Ben Goertzel]: I don't think so. I think the brain is a mess. I think when we see and fully understand it, we will say, goddamn, I can't believe this thing actually let's us think.

[Steve Grand]: Well I disagree with that. I think only *parts* of the brain are a mess.

[Mike Vasser]: Does the problem look better now than it did 20 years ago?

[Sam Adams]: I think the empowerment of the individual scientist is so much greater today than it used to be. Used to be you'd have to have a good budget to build the big system. Now you don't. The tools we have are so powerful, and the components that we have, that a single person can sit down and build a really significantly large system. And so you have the opportunity -- it is sort of like what Steve Grand was saying, this is a problem that is best addressed today at least by having it all in one head because it is just too complicated to split up. And I think that's a real advantage we have now that we didn't have before. Just empowering of the individual mind so that one person actually can build it, or at least build a significant chunk of the system on their own. Drawing from everybody else's research and knowledge, of course, but they can implement it, they don't have a big barrier to implementation.

[Eric Baum]: I think there have been advances in knowledge over the years. It's not like we haven't learned anything. I mean the whole field of learning made a lot of progress. If we go back -- over my career, you go to 1980 and you didn't understand what learning was, and now I think we have a pretty good understanding.

[Ben Goertzel]: I would echo that sentiment. The Novamente system in particular builds on a lot of insights made over the past 20 years or so. I mean the stuff that we are doing in the evolutionary learning module builds on the Bayesian Optimization Algorithm and so forth. And the probabilistic reasoning system we use is unique -- but it would not have been nearly as easy to construct if not for all of the previous work in various aspects of probability theory. There has been a lot of probabilistic reasoning work in the

mainstream AI community. So even though I think the traditional AI community has not focusing its efforts ideally, because of ignoring AGI, by and large, they have done plenty of useful work on various component technologies that can be used to make AGI systems.

[Stan Franklin]: Can I add to that a little? In my opening talk here, I had a slide listing a large number of theories from cognitive physiology that had been embodied in the LIDA system and that couldn't have been done some few years ago, some 20 years ago. It certainly seems to me that when I first started looking at this thing about 15 years, it seemed so far in the future that I thought it silly, but let's do it anyway. Now it seems much more palpable than it did because I have got this cognitive theory of everything that looks as if, if you give it a chance and let it move along, it might in fact be developed into an AGI.

[Eric Baum]: To go back to your question before about the science, the one thing you said was that nothing is repeated and I think that's just false lots of things get repeated – I mean I learned here that somebody repeated Hayek [a learning system developed by Eric Baum, which was replicated by Moshe Looks], but that's a very minor one. For example chess programs are repeated you know multiple times, or support vector machines. I could give a long list.

[Phil Goetz]: I guess I was thinking of general AI systems, such as when Ben Goertzel talks about Novamente, I've not heard of anyone trying to duplicate what they report.

[Eric Baum]: Novamente, in some ways, it seems to me like the systems of the 60's, which were hard to repeat, and maybe weren't repeated and that was a criticism that used to be made, but I don't think that's true with a lot of the progress that has come about now.

[Hugo de Garis]: I'm no expert in brain scanners, but people tell me they are subject to Moore's Law. Talking to a neuroscientist, I asked, what's state of the art in brain scanner work today? She said millimeters and milliseconds. So you know its true that – well it is true that its subject to Moore's Law. Then you get into the Ray Kurzweil scenario of being able to scan the brain to the point of every synapse. So you can imagine downloading all that information into a kind of future hypercomputer, and they're coming. I talk about Avogadro machines. If you know any chemistry, you know Avogadro's number, the number of atoms you hold in your hand, which is effectively a trillion trillion, so you have all the components and hardware needed. You can imagine huge amounts of information on how the brain is structured. Then you can start doing experiments on structure and trying to figure out how it works. I see this eminently doable in the next several decades.

[Stan Franklin]: I'm concerned about brain scanners being... I think they are fast in time, but the ones that have good spatial resolution, takes seconds. I have serious doubts about how far we are we going to get with those. But I'm not an expert there.

[Steve Grand] At the end of that process, all we'd end up with is another brain and we already know how to make them. We also already know how to take them apart and we still don't understand how they work.

[Audience]: How much do you want General AI, and how afraid are you of it?

[Sam Adams]: I want it a lot, and I'm not afraid of it at all.

[Hugo de Garis]: To me it's a religion, and I'm dead scared.

[Stan Franklin]: My then 16 year old daughter, who had read Bill Joy's article, said to me, daddy these intelligent systems, that's what you do isn't it? Why do you do it? The answer I had to give her at that time, which is still my answer, is that I think we are too far

away to do ourselves any harm yet, and by then, maybe we will be able to figure out somehow not to do ourselves harm.

[Steve Grand]: I'm a member of Sam Adams' religion and I want it. And I'm not scared of it. Intelligence doesn't scare me at all; it's the stupid people who scare me

[Sam Adams]: Artificial stupidity is much more dangerous.

[Ben Goertzel]: My response is a little more like that of Hugo de Garis'. I want it and I'm moderately, though not massively or panicked-ly frightened over it. There is obviously the possibility of very grave dangers, yet there are also a lot of other possible dangers that are out there, from things that don't have all of the good aspects that AGI may have.

[Pei Wang]: If I say I'm scared about it, you should not believe me; otherwise I should not be working on it. I think the same thing is for everyone sitting on this panel.

[Hugo de Garis]: Could we take a vote?

[Phil Goetz]: Audience question, raise your hand if you are scared of AGI?

[About half the audience raises their hands]

[Phil Goetz]: Let me ask a more controversial question. How many of you, if you had built your AI, and you thought it was of human-level intelligence, and you really didn't have any means for controlling it, how many of you would turn it on? Raise your hand if you would turn it on?

[Izabela Goertzel]: Human-level is not that smart. [laughter]

[Phil Goetz]: But that's how I am phrasing the question. So that's one, two, three, that's about 10 hands out of 35, 36 people.

[Karl Pribram]: From my experience, you turn it on, and five seconds after its working, it breaks down. [laughter]

[Mike Ross]: The opportunities for collaboration seem so small. There seem to be only occasional opportunities for working together and testing out things. What does the panel have to say?

[Sam Adams]: There is a big difference between people's code and people's ideas. People tend to not want to use other people's code. Programmers, you know, love writing their own code. It's a problem software has in general, but the thing we do have a possibility for, that I've been kicking around with some of the other people here, is just seriously collaborating consultatively with each other. Let's not go away and not talk to each other for two years, and then come back and see how we proceeded. In this day and age, we ought to be meeting at least every six months, as a large group, at least, with a mailing list, a wiki, who knows what else, keeping us in tune and freely consulting. You know these things are so deep, there is no way in the world that you are going to give away the groceries on one of these systems. These are huge complex systems. I can go talk with Stan's team about LIDA, because we are similar enough, that I could probably say you missed a feedback loop there. That may save them years of work, that may be a dead-end they don't find forever, whatever, that's the kind of collaboration I think we need to have to spirit of. Granted, whoever builds this thing is going to win the Nobel Prize, but really I want to see it work, I mean is there anybody who doesn't want to see it work -- hopefully we can turn it off, right? So if Ben Goertzel does it, great, and he uses a couple of my ideas along the way, fantastic.

[Phil Goetz]: Do we need to cooperate now before those floodgates of money that you spoke about open up?

[Sam Adams]: When it finally gets greedy then, then it's over.

[Stan Franklin]: The first thing, the very thing that Ben Goertzel said, and I have already made arrangements with one or two people to collaborate in the kinds of ways that Ben is talking about and trying with others as well. Let me also say that right now I have members of the research team that's working on LIDA all over -- most of us are in Memphis, one is in Berkeley, one is in New Zealand, one is in China, two are in this area, and they all contribute one way or another. Let me say, by way of invitation, there is plenty of room for others if you want to get on board, please let me know.

[Eric Baum]: I agree that the three of you have similar systems, and also the three of you are complaining, well at least two of you, are complaining that you need a lot more funding to finish them. You might consider merging them and making one open source system. You'll have a lot more manpower and you could finish all these different pieces.

[Sam Adams]: Open source of these technologies is a real opportunity.

[Stan Franklin] That's a possibility.

[Pei Wang]: One problem, or difficulty of cooperation in this field as far as I can see here, is that we are actually not exactly working on the same problem. We are already seeing several very different paradigms in the field of AI. Also there is the problem of terminology. I am fully aware that my working definition of intelligence is quite different from many other people. But I do have the research to follow it up. I believe the same for the others. So my best suggestion is that we try to cooperate to push as far as possible. I will try my best to find overlaps so we can cooperate.

[Ben Goertzel]: I largely agree with what others have said. I do think though that, going back to Mike Ross' original question, there is some fundamental intractability in this particular domain, in terms of communication between different teams. I think it is kind of similar to how it's so hard to understand someone else's software code compared to writing your own. Even if it is well constructed code, it's often still quite painful to read and understand the details of someone else's code. Similarly, if you have three AI systems, and they are kind of similar at the high level, it will still be hard for the creator of each system to understand the other ones. You just make so many small decisions for so many complex combinations of reasons, when you build a system like this; and to communicate all these to someone new takes a really long time. If we opened up the Novamente design to someone, it would take a long time for them to understand the things we've done well enough so they can really intelligently give advise and spot something that's wrong versus right. So I am not as confident as Sam Adams that in a brief period of acquaintance someone could look at the Novamente system, or his system, and say, well you missed a feedback loop. I am afraid the things that we are missing, if we are missing things, are things at a deeper level and would take a much greater acquaintance with the system to figure them out. But having said that, although I do think there is a fundamental issue there -- which is because we are engineering complex systems based on somewhat different approaches -- there is certainly the option for a lot more sharing than this workshop affords. Hypothetically if we choose to take the time -- if each of us took one day to describe the inner workings of our system to each other, then spend a day trying to get a common language, then spend another day describing our system in the common language -- we could get a lot of sharing done. I am not necessarily advocating that -- but if we did that, I am sure we'd have a much higher chance of understanding what each other were doing and making profitable suggestions about what each other were doing. The final point I wanted to make, is something easier than making suggestions about the core logic of each other's

system or fusing each other's systems (as an aside, actually I don't think Sam Adams' and my Novamente are similar enough that we would ever agree to make them the same system), is that there are big opportunities to make useful common tools. The example that we harped on throughout this workshop is the AGI-SIM simulation world. If various teams wanted to make a simulation world, we could agree to collaborate on this -- everyone could put in a little bit of work to help with the annoying problems of making legs work right and dealing with different kinds of vision and so forth. And there are probably many other examples of shared resources that could be helpful; but that particular resource (AGI-SIM) is interesting because -- not only is it open source and could be collaboratively worked on, but it could provide an environment for different people's agents to interact with each other, which is kind of an interesting thing. It'd be cool to have an online playground where AI's could just kind of cruise around; and that actually would be very helpful in attracting funding as well. If there was a public, open source AGI-SIM server, different AI teams could log on and let their AI's play together, and that might help us get some of the funding that we want.

[Stan Franklin]: In spite of all the commonality, the major difference is in how we are going represent data. It seems that any such system has to have a common currency that runs through it; we have certainly had to do that with LIDA. In terms of what Steve Grand said earlier, I think our representation can bridge the gap, I think we are somewhere in between, where we move up to the symbolic and then down. This is still mostly conceptual either going down or going up. We have gone up a little bit in the LIDA model, with natural language understanding, and that kind of thing, but not much.

[Sam Adams]: One interesting opportunity that Ben was talking about is possibly cohabitating a common avatar space to train in. The interesting possibility there for collaboration is having one system train another -- and in training each other, teach each other what they know.

[Ben Goertzel]: New possibilities for artificial stupidity. [laughter]

[Ben Goertzel]: Well I feel you have all heard more than enough of my capability to vocalize, so I'm not going to make very lengthy closing remarks here. In general, I'm pretty happy with the way this workshop has gone, I think we have gotten a bunch of really good people, a lot of interesting approaches to the general intelligence problem -- and as many others have echoed, it seems to be close to the time when AGI research is going to take off, in terms ideas and systems maturing, and in terms of the community of researchers growing, more and more new people coming into it. I have been thinking about this stuff for a long time, and working on it for a moderately long time. Others have been working on it significantly longer than me. I'm 39 and I can say that it's definitely been an AGI Winter throughout most of my 20's and 30's. I can see now that it's starting to thaw, in that we can have a workshop like this and you can actually get a bunch of people to come to it. In the last few years, AAAI and IEEE and IJCNN and so forth are having workshops and symposia now and then explicitly touching on human level AI. You can go and give a talk on an integrated AI system or a would-be AGI system and only half of the room points their finger and laughs. So, its really significant progress; and I think the community will grow and we'll see more and more interesting stuff happen. Regardless of whether my own personal optimistic time estimates come true, I feel like the exponential growth curve of AI is definitely cruising upward in a noticeable way, and its certainly going to continue to do so, which is pretty exciting. We don't have specific plans for our next workshop but in

some form or another there will be more of these and we hope to draw in an even broader community. Thank you all for coming and for all your contributions!