

CHRIS FITZPATRICK TALKS WITH FRANCIS HEYLIGHEN

CHRIS FITZPATRICK

I understand your work here at the Global Brain Institute (GBI) to be centered on the idea that the entire world population and all of its knowledge is mixing with machines and the Internet to produce different forms of intelligence, and that this is resulting in a massive network you call the Global Brain, right?

FRANCIS HEYLIGHEN

Yes.

CF

And you're now working with the Vrije Universiteit here in Brussels, but are you already working outside of academia in some ways? Are you working with governments or businesses?

FH

At the moment not yet. Our purpose is basically academic. Although our sponsor is a private company, it is sponsoring us more out of a general interest in scientific research and not so much in terms of direct applications. If those applications came out, they would be the first to be able to profit from them, but that's not our leading intention. The intention at this moment is basically to understand the process by which the Internet is becoming more intelligent.

This is what I call distributed intelligence. Distributed in two senses—it's distributed over lots of different people like collective intelligence, but it's also

distributed over people and machines and technologies. Different components of completely different kinds—like software, hardware, sensors, robots, people—all participate as components of this intelligence; they all contribute with their own specific skills, their own specific information, and their own specific abilities to solve problems or to attack the challenges. And what you need is to coordinate all these different activities. You have an enormous amount of knowledge available in this world, but usually one doesn't know what the other knows, therefore if you have a problem you don't necessarily know where to find a solution for it. What the Global Brain should do is to make all the knowledge available at once, so that whenever you have any kind of question or problem it immediately would go to those machines or people or databases that are most likely to help you to get an answer.

CF

Right. And you've written that over time the Global Brain is becoming more and more responsive. Do you think that it's in some way actually sentient?

FH

It's a question that I regularly get asked, but it's very difficult to answer because people don't have a good definition of what sentience is, or consciousness, which is an even stronger word. It's sentient or conscious in a sense that it senses

what's happening or what needs to be done. Whether it has also some kind of a higher-level self-consciousness?—Maybe not yet at this stage.

CF

It hasn't gone self-aware yet.

FH

No, but that is possible. I could imagine it happening in, say, a decade or so. In a sense what we are doing is helping to contribute to the self-awareness, because at the moment the process that we are discovering is a spontaneous process. The process is not controlled or steered by anybody. Lots of different people, lots of different companies are all contributing little building blocks to the process, but nobody knows what the building will look like. The building is self-organized out of all these building blocks. What we are trying to understand is this process of self-organization, what drives it, and in which direction it is likely to go. By doing that you become more conscious of what's going on and you could theoretically make the process go more efficiently. In that sense our mathematical models could become some kind of self-consciousness for the global being.

CF

It's a reality now, not a metaphor. It's developing at a rapid rate, but since you can sort of steer it or push it forward—in terms of practicality and material, but also along a timeline—are there ways that you could actually make it develop more exponentially? Or is spontaneity what's driving it?

FH

It is spontaneously improving itself, but there are lots of things that are holding it back. Most of these things are not particularly sexy things. These are things

like incompatibility between different kinds of software systems, when people don't really understand how to use certain devices, all kinds of parasitic things like advertisements or memes spreading with all kinds of rumors, all kinds of silly uses of the internet, all kinds of things that eat up a lot of energy and basically dissipate it. And part of what I would call friction means that these are things that are not working well and are dissipating energy just because they haven't been coordinated yet.

Some of these things I would call parasitic modes. This means things that profit from this efficient network for themselves but without doing anything in return like, let's say, viral marketing. What you are doing there is you create kind of a mind virus that you are spreading, not to help the people, but just to help your own company. So that's another form of the dissipation of energy.

And at this moment I think at least ninety percent of the energy that people spend on the net is being dissipated on things that don't need to be there.

CF

Are there people who impose willful obstacles—people who want to complicate and set up roadblocks to disable the Global Brain?

FH

I don't think so. There are people who might like to do that theoretically, but these are typically the people who don't understand enough to know how they would do it. Controlling the Global Brain in that traditional way by having rules and regulations and policing doesn't really work because the Internet itself is just so difficult to grasp, so quick to self-organize and to find ways around, that all attempts to control it using these methods are not going to work.

They may work for very specific things, for example if you want to stop the spreading of child porn you can use some of these methods, for example, report a website to the police, who will then shut it down. That will work for very specific things, but not for this kind of self-organization of the Global Brain. I don't think you can stop that in this way.

CF

So at this point it's already progressed beyond return.

FH

It's progressing because whenever something can be done in a better way, and somebody tries it in this better way, then the better way will be kept and a less good way will be forgotten. And that's a continuous improvement. It's distributed, it's spread over millions of people, over millions of websites, and you can't really grasp that process.

CF

I've read about the Global Brain's reinforcement of its own strengths and the diminishing of its own weaknesses as a part of its progression. In another text you mentioned how space travel has introduced a new paradigm—it has showed us how much more interconnected we are and how we're much less individual than we tend to think of ourselves. Since you're working on mathematical models and simulations of the Global Brain, might this provide a visualization of the Global Brain?

FH

Not exactly a visualization, but a mathematical model.

CF

Could this offer a new paradigm shift in the same way that seeing earth from outer

space has? I mean, is it a way to see the Global Brain through mathematics?

FH

Yes, maybe it is, but...because a mathematical model is very complicated we will need to find a way to kind of visualize the results.

At this moment, I must say, we haven't yet thought too much about it. We first wanted to make sure that we had the right assumptions and the right techniques. We are still at that stage of making sure that it does what it is supposed to do, and then we will try all kind of developments to get to a more intuitive grasp of it. I think at this moment we are still probably the only ones that really have the ambition and the basis to make a mathematical model of this phenomenon. It is very ambitious, so I don't guarantee that the model will be really useful, but at least it will try to bring all those things together.

CF

And have you found that most of the public can understand this with a little bit of explanation, or...?

FH

Well I think they have understood the basic metaphor that the Internet can be something like a nervous system for humanity, but from there on it becomes much more difficult for them to connect it, let's say, to the everyday activities on the Internet. You can see, yeah ok, that it's something smart you do, but there are millions of small things you can do on the Internet that seem to have nothing to do with one another. So the problem is you need a kind of higher level, a more abstract paradigm to see how all these different things fit together. How watching videos on YouTube works together with looking for craters on Mars, how that works together with writing articles for

Wikipedia, how that works together with sending posts on Facebook so that your friends see them. All these things are kind of instantiations of this process. But you need to have this pretty abstract graph of what this process is. We call this process challenge propagation.

A challenge is a problem or opportunity. That means something that challenges you to act, that challenges you to do something, because by doing something you can improve your situation with respect to how the situation was before. If it's a problem and if you don't do something, you'll get some kind of a punishment, you will suffer the disadvantages. If there's an opportunity and you don't do something, you will not get the benefits that you could have expected.

In both cases it's the doing something that gives you benefits. It is better than not doing anything.

So all of the situations that you deal with in real life all around the Internet are in this sense challenges. They are things that offer you some possibility that in doing something you would improve your situation.

And the propagation aspect is that most of these things are not just challenges for you, they are much bigger than you. They offer benefits to lots of people, and if you have gotten out what is good for you, you will tend to pass it on to others who can do more than you could do with it, who can expect additional benefit out of it.

That's the propagation of challenges. And that is what the Internet is so good at. You get an email that suggests you should do something, or that tells you about the problem. "Ah, yeah that's interesting, I have learnt something new, now I know that the next time I should do this or that." And then you forward it to your friends. That's a simple example. There are all kinds of more complex examples, but

the principle is the same—you propagate challenges, and because of the Internet it has become almost painless, almost costless to propagate. The problem is that you need to propagate to the right place, because if you just send it to everybody you know you just contribute to the information explosion. That's not a solution. So the intelligence of the Internet will be in finding out which challenges should be sent to which agents. An agent can be a person or it can be some kind of software or hardware.

CF

So these are the algorithms that could determine the efficiency?

FH

They are the algorithms that could determine the efficiency, but in the end it will be a process of trial and error. You send something to somebody. He likes it—next time he will pay more attention to what you have sent to him. He doesn't like it—he cuts his connection to you or he unsubscribes from your mailing lists or whatever.

CF

So that's the connection right there—of the human and the machine and the network through the tool of the Internet?

FH

At the latest stage we would want to open an automated process, but at this moment it's mostly just a question of trial and error. Things happen and they have the effects, and if they are successful they get reinforced. If they are negative, there are elements to suppress them.

CF

Is the ultimate goal to fuse this human software and all of these machine capabilities in order to create a new level

of intelligence, resulting in new solutions otherwise impossible to achieve?

FH

That's not the ultimate goal, but the general direction is one of reducing friction. Friction is the process of wasting energy on needless transactions and activities. In order to communicate with somebody else, let's say in the old days, you first had to write down your thoughts by hand in a letter, and then put it in the envelope, and then put a stamp on it and then give it to a mailman and so on and so on. There was a lot of wasted energy along that process. But now, if you type it in an email and you just push the "send" button, all the rest goes automatically. All that friction has been eliminated. So in that sense—typically in the interfaces between humans and machines—there is still quite a lot of friction; you still need effort to talk to a machine because a machine doesn't understand what you are saying. That is the kind of friction that will get eliminated, and has already been very much reduced—but there's still a lot of space there for further reduction of that friction. And if there's no more friction, it would be as if in having thoughts they immediately get into the machine and the machine answers. At that moment you can say that the machine has become a part of your brain. There is no longer a separation.

CF

One of the things I find really interesting about trying to imagine the Global Brain as an image, or as one interconnected thing, is the expansive sense of corporeality it brings to mind. Of course, the Internet is a massive space and it's quite immaterial, but it's run on servers which need to be fueled, and which need to be monitored, serviced, etc. There's a hidden corporeality. Would you say that the Global Brain is functioning similarly?

FH

Yeah, for me the Global Brain...

I always put it in conjunction with the global super-organism—that means the brain has a body; the brain is not just a disembodied thing. It has a body and a body is all the parts that interact with the outside world. And these are people but also machines...all machinery. If, for example, you would have some kind of Global Brain system that would control the traffic lights in the city, determining which trucks pass at what time where, it would be actually interacting with the material world, and we are going toward this. This is what some call 'industrial Internet' or the 'Internet of things', where tiny sensors are connected to all kinds of objects, machines, books, things that are stored in warehouses, and you can control which thing is going where; you kind of make this direct connection to the physical world and the virtual world of the information. This will happen more and more.

CF

That's incredible. You have, however, also written about the potential negative side effects of network dynamics.

FH

What you have in all processes of self-organization—meaning in this spontaneous development—is a danger of positive effects getting out of control. Self-reinforcing processes that become like vicious circles which produce more and more of the same thing, even after you have actually already reached the level that you wanted. That is what you need to be careful about. Financial crises are a good example—in speculation, everybody buys more of the stock because they expect that the price will go higher, so it goes higher because everybody buys, up to the point where somebody starts to say, "yeah, but

the price actually is much higher than the thing is worth in the world, I better stop it now.” The whole thing then goes on but in the opposite direction—the price goes down and goes down and goes down. This is called “network effect”, but I prefer to call it positive feedback—something that reinforces itself via a number of loops. This is an intrinsic danger in networks. You can’t protect against this but you have to be aware of what the dangers are.

CF

In a recent talk at Objectif Exhibitions in Antwerp, Adam Kleinman spoke about the Black-Scholes formula, which Nobel laureates developed to try to predict the stock market. This is all an oversimplification of course, but what I found interesting is that they didn’t factor in the possibility that someone could just say, “I’m not going to pay.” I cite this example because something perfectly automated can always fall apart through human error, or be circumnavigated by human decision-making. In this way, are we humans also a threat to the Global Brain?

FH

Well, normally in the Global Brain human error should not produce much of the problem precisely because it’s the idea of the collective intelligence—you aggregate the opinions of lots of people. So if some of those are wrong, and others are wrong but maybe in the opposite direction, then in aggregation, they tend to be more balanced even though there is no guarantee of that. Again, what you should be careful of is positive feedback—if different people influence each other, then the opinion of the few can become opinion of many, and can become an opinion of everybody.

So there is always this danger of infection; of the self-reinforcing influence of the one on the other. That is what you

should be careful of. But it depends on the kind of medium. In the web it’s not too dangerous, in the sense that you have an enormous diversity of websites. Some get more hits than others, but there’s always enough of a diversity to create a balance. But you might get the problem where people only start following the news that they themselves like, and that others like them like, and those people would in the end get a very selective and biased view of how things are, just because the process reinforces itself. That is dangerous.

CF

That reminds me of when people post things on Facebook to their own closed network of likeminded “friends.” A network of similarity also becomes a barrier.

FH

There is a well-known problem in social psychology called ‘groupthink’. People in a group start to so much reinforce each other’s opinion that they develop a kind of conviction that they are right, but at the same time they are completely disconnected from the reality. You often see that in cults, sects, also in extremist groups like Al Qaeda. Since the one constantly confirms what the other says, they get more and more convinced that they are the ones who have the true view of reality, and tend to dismiss anything which comes from outside of the group. In the Internet there is a danger for such things; but this danger has always existed. Let’s say, it’s easy to create closed communities. It’s easy to open them up on the Internet, but you have to want to open them up. And those who don’t want to open up can create their own closed communities. In this way they reinforce their own opinions, so much so that they get completely out of touch with reality.

CF

And that's a clot in the stream of the Global Brain?

FH

If you look at it in terms of organs it might be something like a tumor. A tumor is a group of cells which grows without regard for the larger organism. It just grows for its own sake. And yes, somehow we will need to develop some controls for that... I'm not quite sure yet how to do that. I think it's the biggest danger that is not yet resolved. It's not intrinsic to the Global Brain or to the Internet; it exists as it has always existed, but the Internet facilitates everything—it speeds up everything, including these kinds of processes. In Belgium they speak now a lot about the radicalization of young Muslims who go to fight in Syria after some crazy ideal that probably was propagated via small groups. These small groups are probably using Twitter or Facebook or email to keep in contact with each other. Let's say that before the Internet it would have taken them much more time to organize something like these organized excursions to go to fight in Syria. Things like these would have taken maybe a decade to organize whereas now they can be done in few months.

CF

So it's about who is holding the tool, no matter which tool. If we consider 'the physical world' and 'the Internet' to be two distinctive worlds for a second, it seems almost as if all the same agents and all of the same problems and situations that are out here are the same within the computer, more or less. But other than its orbiting satellites, the servers, machines, and software running the Internet exist here on earth. In the future, will there be any distinction between the two anymore? Or will there always be a distinction?

FH

No, I think that the Internet and the physical world will merge. Everything in the physical world will communicate with the Internet. And everything you do on the Internet sooner or later will have some consequences in the physical world. We will need to develop a number of safeguards for problems that most people at this moment are not yet aware of, or how serious they could potentially become. For example, I think at the moment lots of people are worrying about privacy, but I think privacy is less of an issue than this kind of groupthink. I mean, you can set a number of safeguards, you can have rules that certain kinds of data cannot be used by certain actors without explicit permission, and so on. There exists cryptography—all kinds of methods exist to protect data, and there exists some degree of law about the protection of privacy. Of course these are still not sufficiently developed for the world of the Internet, but people are aware of the issue and they already have some tools to deal with it. The kinds of explosive revolutions that spread via the Internet can be very positive, like the revolution in Egypt, but they can be negative, like the riots in London. We don't know yet what the worst of its effects may be. Another example I discussed in one of the papers was the genocide in Rwanda—before the days of the Internet—which was spread in part by a radio station. Radio is kind of like the precursor to the Internet. It's a mass medium. A radio station was used to distribute propaganda and to set up the Hutus against the Tutsis. Historians say that these radio stations played quite a nasty role in the genocide. Now imagine something similar happening via the Internet. It's quite possible.

CF

So in a way, the groundwork you are laying now, on the understanding you are gaining, is a protective tool for the future.

FH

Yes, because if you make a mathematical model in principle you can simulate these kinds of things. People have already been making smaller-scale simulations of what they call social contagiousness, where a fashion, a rumor, or a false idea is spread from one person to another to another, in order to see—depending on the structure of the social network and some other parameters—how far it spreads. What these kinds of models do is give you a certain degree of prediction. An example I once mentioned in a discussion with my colleagues: suppose that you are in an explosive situation, like in Iraq or in India, and there is a rumor somewhere that the Muslims in a particular town have destroyed a Hindu temple. The rumor is wrong, but you can predict it will spread very quickly. If you know the social network—you have for example the Facebook network of the people in that region—and you can guess that the most extreme Hindus will be the first to spread that rumor toward their friends, and their friends will spread it to some of their friends and so on, then you can kind of anticipate what will happen. Then what you could do is inject a counter-message to exactly those places where you expect the false rumor to appear. You could have a kind of immune system that tells them, “You may hear the rumor that Muslims have destroyed the temple, the fact is that nothing of the kind has happened, but maybe there was a minor incident there or there.” You could do that. That is what these models could potentially do.

CF

It’s amazing to think that you are building an immune system then, when you put it in those terms.

FH

Yeah, that will be one of the applications. First we will need to establish the model that works. And once we have it, then we can use it for these kinds of applications.

