Distributing Cognition:
from Local Brains to the Global Brain

(v3.1)

Clément Vidal
Center Leo Apostel
Global Brain Institute
Evolution, Complexity and Cognition research group
Vrije Universiteit Brussel (Free University of Brussels)
Krijgskundestraat 33, 1160 Brussels, Belgium
Phone +32-2-640 67 37 | Fax +32-2-6440744
http://clement.vidal.philosophons.com
clement.vidal@philosophons.com

To appear in: Goertzel B. and Goertzel T. (eds): *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*

Abstract:

We show how the externalization of our local brain functions is leading to a planetary level intelligence, or global brain. We argue that this mechanism of externalizing cognitive functions is a fundamental driver towards an ever smarter human-machine symbiosis. We discuss implications and applications of externalizing various cognitive functions such as memory, computation, hearing, vision, brainstorming, reasoning, navigation, emotions and actions. We illustrate the scenario with a fictional story of a day in year 2060. We then take a top-down approach, and argue that this progressive externalization helps to better understand, foresee and facilitate the emergence of a globally distributed intelligence, best conceptualized as a *global brain*. We discuss possible symbioses between biology and machines, and what would be the elementary elements composing the global brain. We finally embrace a broader cosmic perspective and argue that even if the singularity is near, an energetic bottleneck is nearer. We suggest that other extraterrestrial global brains in the universe might have successfully dealt with this energetic bottleneck.

A distinctive feature of the human species is its ability to create, use and refine tools. Doing so, we dramatically extend and enhance the variety and power of our senses, actions and information processing (see e.g. Clark 2008). The externalization of cognitive functions has produced major transitions in the history of culture. For example, the externalization of memory is the invention of writing, while the externalization of calculus is the invention of computing machines.

Can we foresee other externalizations and their impact? What will the ongoing externalization of cognitive functions ultimately bring? Externalizing cognition first empowers the self, but also disruptively change society in a second stage, as the examples of writing or computers clearly show.

My goal is twofold. First, to use the externalization of cognition as a bottom-up framework to systematically anticipate and explore potential new technologies. Second, to discuss top-down approaches to the idea of an emerging globally distributed intelligence, or “global brain” (see e.g. Heylighen 2007a).

The global brain vision has its roots in organic conceptions of societies, and is in fact focused on the analogy between a nervous system and information processing in our global society (see Heylighen 2011a for the history of the concept). The focus of this vision is thus on *information* and cognition. But the distribution of *matter* and *energy* are equally important to study – although not the main focus of this paper. For example, we could imagine a global distribution of matter, or "matternet". It would operate a global distribution of raw materials, while tridimensional printers (see e.g. Lipson and Kurman 2013) would fabricate any kind of object locally. In a similar vein, the global distribution of energy, or the emergence of an "enernet" would lead energy consumers to become also local energy producers. Energy thus becomes distributed in a bottom-up instead of a top-down fashion. From a systems point of view, I just would like to remind my readers that a "matternet" or an "enernet" are as important to take into consideration as the internet to understand and anticipate the future of our society.

To characterize internet intelligence, a systems theoretic approach is also required. We can distinguish four stages of information processing (Heylighen 2013): input, processing, output and feedback. For a system to be intelligent, it must perform these stages as well as possible. The stages can be summarized as follows:

1) *input* is intelligence as information collection
2) *processing* is intelligence as interpretation and problem solving
3) *output* is intelligence as efficient and relevant action
4) *feedback* is intelligence as capacity to learn

An intelligent agent is able to collect relevant information (step 1). Generally, the more information, the better. However, it should also be able to process it in real time, or to store it. If you are given a book printed with the raw data of the human genome, there is little to do with it. So, processing information is necessary to make sense of the information gathered (step 2), and to use it for action (step 3). Finally, feedback is essential to learn (step 4) and therefore to fine-tune inputs, processing and outputs. Heylighen adds that each of "these functions can be performed in principle either in a distributed or in a centralized manner." This distributed versus centralized distinction is important to keep in mind in the context of our information society, where the distributed manner becomes more and more prevalent.

In the first section, we describe how the externalization of cognition brings new levels of intelligence. In section 2, we explore a fictional story of what a day would be like in 2060, with externalized technologies. In Section 3, we step back and discuss the global brain vision. We ask what elementary parts will constitute the global brain. Humans? Machines? A symbiosis between the two? A more complex ecosystem? Finally, we take a cosmic perspective in Section 4, and show that if the singularity is near (Kurzweil 2005), an energetic bottleneck is even nearer! We speculate that extraterrestrial civilizations might have already overcome this bottleneck. The combination of both bottom-up and top-down analyses is insightful because it connects the dots between our present, our near future towards a global brain, and our far-future beyond the technological singularity.

1 - Local Brains or the externalized parts

Most tools derive from the externalization of our bodily functions. They concern primarily matter-energy processes. For example, wetsuits or spacesuits are enhancements of our skin; horses, cars, trains or planes are locomotion systems which extent our motor functions. Here, we focus on information processing technologies which enhance cognitive functions. As Dror and Harnad (2008, 3) argued, cognitive technology does "extend the scope and power of cognition, exactly as sensory and motor technology extends the scope and power of the bodily senses and movement." They add that just "as we can see further with telescopes, move faster with cars, and do more with laser microsurgery than we can do with just our unaided hands and heads, so we can think faster and further, and do more, with language, books, calculators, computers, the web, algorithms, software agents, plus whatever is in the heads of other cognizers."

This approach emphasizes more *Intelligence Augmentation* (IA), than *Artificial Intelligence* (AI) (Fischer 2006). As we will see, this IA perspective naturally accounts for present Information and Communication Technologies (ICTs), and promises to anticipate future techno-sociological development. In the cognitive science landscape, we are thus naturally in line with the tradition of situated and embodied cognition (see e.g. Wilson 2002; Anderson 2003; Clark 2008).

In the following analysis, we distinguish three phases of the externalization process. First, the *bodily self* with no extension; then the *extended self*, using tools and the environment to enhance its capabilities, and finally the *networked self* when it is connected to the internet. The *bodily self* performs information processing on a very local scale. It collects internal or local inputs, does centralized internal processing, produces local outputs and relies on local changes to tie his feedback loops.

The *extended self* is much smarter. It extends its inputs with the help of tools, such as goggles, its processing by using other tools such as computing machines, its output by using tools such as hammers and shoves, language, pen and paper. The feedback loops operate through these extensions, which can be repaired and enhanced. For example, an author will do her best to repair or replace her failing favorite pen or computer, which are almost part of herself. Language itself can already be analyzed as an extension of communication, operating not anymore on the individual bodily level, with hormones or neurotransmitters, but across a social group. This social aspect is essential because cognitive tools can be shared. For example, thanks to teachers, our social partners, we learn at school how to speak, write and calculate. Yet, input, processing, output and feedback remain local, in the sense that it is not possible to listen or speak to thousand people at the same time, nor to learn from the experience of millions of people.

The *networked self* is much smarter thanks to its connection to the internet. The scale of its input, processing, output and feedback is global. It can communicate with any other networked self on the planet. The benefit stems not only through connecting human brains together, but also connecting humans to machines, i.e. *artificial agents,* which can be software or hardware agents. For example, internet users rely daily on smart algorithms and big databases to search or filter the available information on the web. I will argue that the symbiosis between social interactions and artificial agents provides and will continue to provide the most useful and disruptive technologies. The extension of cognition started as an *individual* intellectual empowerment (e.g. with pen and paper). But such an extension also has a *sociological* impact, since it opens the way to share tools and information with social partners, resulting in faster and wider communication on a social scale. For example, the invention of printing lead to the massive distribution of information via books and newspapers. In the end, we interpret the singularity as a symbiosis between humans and machines leading to a major evolutionary transition, the emergence of a globally distributed intelligence, or *global brain* (see Section 3).

Let us now explore more specific implications of externalizing various cognitive functions such as memory, computation, hearing, vision, brainstorming, reasoning, navigation, emotions and actions. Importantly, since memory, bandwidth and processing increase superexponentially (Nagy et al. 2011), we should not consider them as bottlenecks in imagining future technological applications.

Memory is implemented biologically with neural networks. It became sharable with language, and has been radically extended and safeguarded thanks to the invention of writing. This allows sharing of manuscripts, first locally, then more broadly with the invention of printing. For a detailed account of the disruptive nature of the inventions of language, writing and printing, see (Robertson 1998). The third phase of the externalization is the distribution of memory over the internet. The invention of hypertext, and later of the web is an improvement on the invention of writing, and can be analyzed as a globally distributed, collective and dynamical memory. To invent the web, Tim Berners-Lee (1999, 4 and 41) was indeed inspired by the functioning of the brain, which can link any arbitrary piece of information. He wanted computers to be able to do the same.

Let us illustrate this point with the collaborative Wikipedia encyclopedia. First, a shared memory obviously accelerates human-human interactions because of the availability of shared information in the environment. Wikipedia editors share and improve constantly the external memory via a *stigmergic* process (Heylighen 2007b). Second, human-machine symbiosis is possible when humans enter data and metadata in a semi-structured or structured way. Such data can then be used to query relationships and properties, including links to other related datasets. More complex or extensive inferences and queries can be processed thanks to software agents, which will increase human knowledge. In turn, humans can use software agents to improve the datasets and establish a human-machine symbiosis. An example of software agents are Wikipedia bots. More than 1600 of them operating as you read these lines. Bots most often leave tags or flags on wikipedia pages, instead of directly and aggressively editing a page. Such a mild contribution is still very helpful, since it can stimulate or nudge humans to work on specific pages such as pages having a very active “talk” page, or which are running an “edit war” (see e.g. Kittur et al. 2007). However, machine-machine interaction is also possible, if bots become intelligent enough to make reliable edits, further edited by other bots.

Computation, or counting and calculating probably started by interacting with the environment. For example, if a farmer wants to count its small livestock, it is easier to establish a one-to-one correspondence between stones and animals than to count mentally. This correspondence procedure is actually the modern definition of cardinal numbers in set theory. From a situated and embodied cognition viewpoint, this means that environmental scaffoldings to cognize (such as stones) can precede their integration into the bodily self. Combined with the inventions of writing, calculus took another leap. Human computation is severely limited by the working memory capacity of the brain. This is why mathematicians extend their cognition with their favorite tools: paper and pen. Later, the specification of the general purpose computer stemming from the work of Church, Gödel, Kleene and Turing led to the general computing devices we use daily. The extension of computation to the internet leads to distributed computation and grid computing.

The hearing function starts with the biological ear, is captured thanks to microphones on physical supports such as tapes, compact discs or memory cards. Thanks to big data and pattern recognition algorithms, automatic music identification services have emerged. They are portable thanks to their integration to smartphones connected to the internet. Their accuracy in recognizing songs is already quite impressive. However, with a hybrid approach of combining algorithms and crowdsourcing to humans, the music recognition could give even better results. For example, humans would be able to guess live music which has never been recorded, where algorithms would typically fail.

Vision is biologically implemented with the eye, but visual information can be stored thanks to cameras, technologically extended to scales with microscopes and telescopes, and extended from the visible to the whole of the electromagnetic spectrum. When vision is networked from smartphones to crowdworkers, such a connection allows new applications. For example, it is possible to get nearly real-time crowdsourced answers to visual questions, which is extremely helpful for visually impaired persons (Bigham et al. 2010). Note that such pattern-recognition tasks are offloaded to other humans, but they could also be offloaded to artificial agents. For example, some search engines allow to search not only by keywords, but by pictures freshly taken. Finally, a mixed approach could prove even smarter. It would first query artificial agents, and then, if they fail or if their confidence in the results is low, crowdwork would be requested. This would be most efficient because mobilizing humans for a simple task that an artificial agent could perform would be a waste of human workforce.

The faculty of *divergent thinking* is essential to cultivate creativity. If we are unaided by external cognitive scaffoldings, divergent thinking simply means having a lot of ideas. But most of the time, we have a lot of ideas thanks to triggering lists or brainstorming sessions with cognitive partners. And of course, we don't want to lose such ideas, so collection tools such as notes, mind maps, or voice recordings are essential in the process. With the internet and the web, it is possible to radically scale up this creative process. Indeed, we can now seek input from a large number of people (e.g. with wikis or with social media websites). The difficulties shift, namely to the aggregation and sorting of the results to make a coherent whole. There are a few collective outlining or mind mapping softwares, but they are probably underused given the benefits they could provide. Software agents could also help divergent thinking, for example by automatically aggregating similar ideas, or even by systematically suggesting combinations of existing ideas. For example, a software could simply take inputs from humans, such as idea A and idea B, and ask "What would happen if you would combine A with B?"

The faculty of *convergent thinking* enables humans to produce arguments which solve problems. It goes hand in hand with divergent thinking, but this time the goal is not variation of ideas, but selection of ideas, to solve a particular problem. Extending reasoning to the environment leads to the use of logical diagrams (e.g. with the theory of constraints of Goldratt and Cox 1984), thus providing greater control over problem solving and argumentation. The reasoning then becomes easier to criticize, revise, understand and discuss. Such techniques are widely used in finance, distribution, project management, people management, strategy, sales and marketing. Distributing reasoning on the internet is still in its infancy, but holds great promises for the future of distributed governance, since it has the potential to lead to large-scale decision making (see e.g. the projects of Baldwin and Price 2008; as well as Iandoli, Klein, and Zollo 2007). Artificial software agents could help semi-automatic reasoning by helping users and groups to systematically and critically question an argumentation. More sophisticated agents could also extract reasonings from existing text, or even reason on their own (e.g. with logical programming). Again, a mixed approach of computers suggesting inferences, and asking for feedback from humans seems more practically useful than pure artificial agents whose reasoning would be limited by the structured data it can gather. It could lead to a real symbiosis, where algorithms would learn from the feedback of humans, and humans would benefit from logical connections suggested by software agents that they might otherwise have missed.

Navigation is hard to operate in an unknown environment. Maps make navigation much easier. Helped with machines and a Global Positioning System using a voice synthesizer, the task of navigation becomes incredibly simple. The connection to the internet can provide additional real-time information, can allow users to correct maps when they find errors, and leads to an increasing variety of location-based services and applications.

Emotions play an important role in cognition. The subject matter of affective science is to study motives, attitudes, moods, emotions, and their effect in personal and collective behavior. Expressing emotions started thanks to body and facial language. It then extended to sounds and articulated language. Emotionally intelligent people are aware and tuned-in with their emotions. They are able to communicate them in a constructive and non-violent way (see e.g. Rosenberg 2003). Emotional intelligence also include empathy, or the ability to recognize emotions in others and to react appropriately. Cognitive reappraisal is a technique to regulate emotions, where the patient uses another person's insight to change perspective on a negative situation. Morris and Picard (2012) have shown the feasibility of crowdsourcing cognitive reappraisals. Other applications could include the crowdsourcing of empathy, which would be most helpful for autistic people. Of course, in many such applications of distributing cognition to the internet, the rapidity of the response is critical to make the tool valuable. If it takes five minutes for an autistic person to take a picture of an other person, send it to a server of crowdworkers, wait for a response, wait for the result to realize that the person was very angry, it would be socially useless. But if the autistic person wears augmented reality goggles or contact lenses, which recognize in real-time faces and their emotions to feed them back as explicit signals, then the tool becomes extremely helpful. Such a scenario lies in a very near future, since we have already algorithm which recognize not only faces, but also human emotions. As a day-to-day example, modern digital cameras have smile recognition features, to take pictures only when subjects smile. There are many more applications yet to be discovered, but there is no doubt that the field of *affective computing* (see e.g. Picard 1997 for a review of the theory and many applications) holds great promise to develop better human-computer interaction and symbiosis.

Let us further explore future scenarios where affective computing play an disruptive role. A key element to make the human-machine symbiosis functional is to have high bandwidth of interaction. For a smartphone user, taking a picture still takes the time to decide to take the picture, to put one's hand in the pocket, to perform a few swipes, wait for an autofocus to happen, and to click the shutter button. This is extremely quick compared to what photography was in the 1820's, but extremely slow compared to what it could be. Imagine you are wearing augmented reality glasses. You blink your right eye, done, your picture is ready to be shared – if you wish to. Better, you don't even take the decision to take a picture. The computer you wear monitors your emotions in real time, and take pictures or videos automatically as you live your life. Later, you can consult this log, and filter it by the most surprising, enjoyable or disgusting experiences you had. How can your emotions be monitored? This is a technical problem, but emotion sensors can already perform facial and emotional recognition. Other sensors such as brain waves, heart rate, skin connectivity or even real-time blood analysis (Golan et al. 2012) could be used.

Let's take another example of affective feedback applied to software development. You use your favorite software, but can't find the functionalities you need or you experience unwanted bugs. The regular way to complain is to write or phone the technical support, or to register to the relevant forum and post a message. But how many users do actually do it systematically? Now imagine that your webcam records and recognizes your feelings while you are using the software. This data is sent anonymously to the developers along with a log file describing what you were doing. The developer could then analyze this data, why the user was unpleased, and improve the software accordingly.

Now, let us imagine how urban traffic could be improved thanks to affective computing. The costs of road crashes are very high, more than US$517 billion (Jacobs, Aeron-Thomas, and Astrop 2000, 11), without counting the pain, grief and suffering of people involved. Let's imagine that most pedestrians, cyclists, car drivers have emotion sensors, linked to their location information. Suddenly, a car makes a huge acceleration, a lot of noise, without any other reason than showing-off. A dozen of drivers, pedestrians and bikers get stressed out. This emotional feedback is anonymized and could be sent directly to the dangerous driver, or processed by the police. The driver would loose points for his driving license, or collect malusses for his insurance. On a nearby road, a driver nicely anticipates the coming of a mother with her young children on a pedestrian crossing. They smile at each other, the driver pleased to respect the safety of pedestrians, the mother pleased with the behavior of the driver. Remember that the emotions are recorded, and the driver would thus collect bonuses thanks to his cordial behavior. Another pedestrian witnesses the scene, and further confirms this cordial behavior. Now, imagine that such a system is implemented on a large scale. The collection of bonuses could be incentivized in many ways, and would certainly result in changing the behavior of drivers. The collection of malusses could also be sanctioned in many different ways. The rules could be even stricter for self-driving cars, for example requiring them to ensure minimum fuel consumption as well as maximum comfort for passengers, by accelerating and braking smoothly.

Our ability to act can be enhanced by using to-do lists. By writing down next physical actions to do, this process gives us power on our actions because we can then reorganize and reconsider them in a very efficient way (see e.g. Heylighen and Vidal 2008). Collective action can be efficiently implemented with job ticketing systems, as it is practiced in call centers or open-source software development. For example, in open source development, users or developers can create a ticket to signal a bug or request a feature; other users can vote to prioritize jobs. Yet other developers will actually fix the bug or implement the new feature. This process happens in a stigmergic fashion, where there is no central planning or planner. We saw that Wikipedia is a typical example of such stigmergic collaborative work.

Softwares using location information can improve action efficiency by filtering an electronic to do list, for example suggesting you to buy a specific item only when the shop is open, when the item is in stock and when you are near it.

An important development in the distribution of actions is the rise of *microwork*, which consists of tasks completed by many people on the internet. Examples include Mechanical Turk by Amazon, and such globalized work rises new socio-political issues regarding the difference in wages between developed and developing countries.

The distribution of physical action remains a challenge. Surely, there are microsurgerons operating robotic surgery via the internet (Anvari, McKinley, and Stein 2005). But other applications such as telepresence and action at distance remain marginally in use. This is due to the difficulty of building autonomous or semi-autonomous robots with efficient sensory-motor effectors in our real and complex world. Biological organisms are still much more sophisticated and efficient than robots for these real-world tasks.

We have described the externalization of cognitive functions and how they change social interactions. We insisted on the impact of combining this externalization with a global interconnection through the internet. The most interesting applications will likely be at the interface of several externalizations. Let us give two examples. I have been thinking about an application called *augmented conversation*, only to realize after triggering this thought to our global external memory (the web) that such a project was already in development (Chorley 2013). As Chorley writes, it consists of “a scenario where people are able to take part in a conversation and as the conversation progresses relevant content is retrieved from the web and displayed to the participants.” Such features could be enhanced with emotional recognition, for example, searching more actively and thoroughly when participants display surprise or facial expressions indicating disagreement. Emotional tagging could also be helpful for automatic logging of conversation, so that the most impressive, unusual, i.e. informative aspects of the conversation would be highlighted for later reference. As usual, not only artificial agents can be used, but also humans. For example, crowdworkers could be paid to listen to conversations and check the accuracy of participants' statements.

The second example is the interweaving of human and computer computation in order to translate quickly, cheaply and reliably. Minder and Bernstein (2012) designed a framework to make the most of human and computer intelligence for this task, and manage to make a good translation of a whole book from German to English in only 4 hours and for about 70$. The process involves first automatic translation of sentences, which is then shown to different humans, who rewrite a better translation. Then other humans vote for the best translations. The process involves quite some redundancy in fixing and verifying the intermediate steps, which leads to very good final result, considering the time and money spent.

2 - A Day in 2060

To see more concretely what the impact of such technologies could be, let us imagine what a day circa 2060 could look like. What follow is a speculative science-fictional story, but I will analyze it and give hints and references explaining why it is plausible.

You're walking on the street and meet a Japanese woman. She starts to speak her native language, but your augmented reality lenses automatically translate and display what she says. Akemi explains that her GPS doesn't function well and asks if she can connect to yours to find her way. You accept the request, but your augmented reality vision also display: “try something”. What happened in the background of your extended self in order to suggest this? In a fraction of a second your sensors and artificial agents did the following. They took a picture of Akemi, from which an image-search was launched, along with face-recognition. Several webpages of her public profiles were found. This information was integrated to create a profile, summarizing her professional, and to a lesser extend, personal interests. Additional visual and olfactory sensors on your wearable clothes did notice unusual pupil dilatation and pheromone concentration. Intellectual and bodily data concluded – on both sides, since Akemi did of course do a similar search – that this encounter was an excellent love match. You could have configured your digital agents to give a you better tip than “try something”, but you chose a low advice specificity profile, to leave some of life's spontaneity. So, you indeed try something, and invite her to join you and your friends for swimming with dolphins this afternoon. You share time and GPS coordinates and you are thrilled that she accepts the invitation.

You run back home, cross the street without even looking at cars. A car brakes violently. You are surprised to see a driver in it, and shout: “poor and dangerous biological human, buy yourself a self-driving car, your reflexes are too slow!” Your emotional reaction was monitored and the automatic legal decision making actually gives you a one bitcoin fine, because you should not have had crossed so quickly the street in the first place, and you should not have had insulted the human driver, which had a negative emotional impact on him. Your augmented reality informs you sympathetically: “I understand that you felt upset and need more security. This security indeed implies that people should switch to stronger human-machine symbiosis, but can you please be more careful next time? This transition towards human-machine symbiosis is still in progress. The driver felt embarrassed and miserable about this situation, which is one of the reason why you had to pay this fine. I don't advice to make an appeal, it will only cost more money and given that the situation was recorded by 10 different nearby sources, there is few ambiguity, so the judgment has a 99.9% confidence. The bloodstream of the driver has also been checked and it was perfectly clean, whereas your adrenaline levels were unusually high.” You understand this but still wonder why human-driving cars are still allowed to circulate. Probably a lobby of the old-fashioned Association for Biological Human Rights.

When you arrive home, a self-driving car just brought fresh cartridges, automatically ordered by your 3D food printer. As soon as you plug the nutrient cartridges in, your 3D printer cooks for you, based on inputs from nanobots floating in your bloodstream, which monitor the nutrients you need most. Your 3D printer is furthermore configured to follow your preferences, in this case, to follow a paleo diet because you decided to be in line with evolution. The animal protein supply is a mix of artificially grown meat, fish, worms, and insect proteins. The food quantity is also higher than usual, because your printer anticipates your sport activity planned in your agenda. Indeed, you gave access to your agenda to your printer. The recipe is a new creation, because you've configured your printer to never print two times the same meal. Life is too short and the world's diversity of cooking too great to eat two times the same meal.

When you arrive at the harbor, safety procedures are quick and simple, just to give your stem-cell box, which could be used by the first-aid-kit on the boat. The boat is small, and no oxygen bottles are taken on board. Instead, the trainer takes a suitcase containing syringes. Just before going into the water, the trainer gives a shot to all participants. What is in the shot? Mechanical artificial red cells, providing a 4 hours in vivo Self-Contained Underwater Breathing Apparatus (SCUBA). You and your friends dive in the water, play and communicate with dolphins, thanks to the dolphin speaker interfaced with your augmented-reality diving mask.

Suddenly, the boat radar displays an alert on your mask: “Shark approaching at high speed; no time to swim back to the boat. Fight is the only option”. But you use your biological brain and think that there must be another way. You remember that dolphins can sometimes fight a shark. You turn to the dolphins hastily, set your dolphin speaker to beam a help signal, along with the 3D shape of a shark you quickly downloaded. Fortunately the dolphins understand your message, they do thank you, but get scared and swim away! The AI advice was wise. You feel frustrated that AI was once again smarter than you.

Out of sea mist, the shape of a shark is coming. Some last minute information is displayed on how to fight a shark to you and your friends. You start to read them, but too late, the shark has chosen to attack you. You see the shark's jaw dramatically expanding and... nothing. You loose consciousness.

You wake up on the boat, fully recovered. Akemi is looking at you. You ask her: “what happened?” She explains that your friends managed to scare the shark by fighting him from multiple sides on its gills, and that he finally released you. You ask: “but how come was I not wounded?” Akemi: “You actually almost died! Your nano health bots detected your right kidney and your liver were critically failing. The messaged was transmitted to the first-aid kit on the boat, and the 3D organ printer started to differentiate your stem cells and printed at fast speed two new organs. I contacted a japanese surgeon expert in organ transfers for an urgent tele-operation. I gave him distant access to the first-aid robotic surgery apparatus, and he could work with the printed organs. I hope you don't mind we chose a human surgeon, we are still not confident enough with the cheaper fully robotic AI surgery.” Your health insurance reckon that the incident could not have been avoided, and financially covers the whole operation. The file is already closed.

You ask: “what about the shark?” Akemi continues: “Since it drunk on your blood, it will be infected by artificial viruses. I guess you feel resented, but you know that global eco-regulations forbid to reprogram them at a distance to kill the shark. However, thanks to this artificial virus infection, the shark is now trackable and should not create any further incident to any diver with an augmented-reality diving mask.” As you put back your augmented reality lenses, you look at your information feed, and see that you have been thanked by diving, surfing and fishing associations for successfully tracking an additional shark.

On the way back to the coast, you skim some news and learn that a bomb has exploded at the headquarters of the Association for Biological Human Rights. The police has found out that the bomb was printed directly through the local 3D printer of the association. The cyber-attack left traces distributed around the globe. Police said the identity of the hacker is uncertain, and the debate rages whether it was triggered by a human or a coalition of artificial agents. At the end of the day, you realize that AI agents have done much for you today, and are in awe and grateful to them and your friends. You owe them all your life.

Let us now analyze this story and give some hints of why it is plausible, if we extrapolate some existing technologies. Augmented reality contact lenses will surely come, and prototypes are being tested (Lingley et al. 2011). In the story, humans are augmented with sensors, noticing details too hard to consciously perceive, such as the amount of pupil dilation or pheromone concentration. Three-dimensional food printers (Cohen et al. 2009) and biological organs printers (e.g. Mironov et al. 2003) already exist in embryonic forms. Automatic translation is supposed to work very well, and could be made more effective thanks to contextual data inputs. For example, your location in the street makes it likely that the conversation will be about finding your way, and the profile of the person restricts the likely vocabulary usage. Machine-machine interaction occurs when your GPS signal and maps are shared with Akemi's system, or when your artificial agents collaborate to tell you “try something”.

Regarding the car incident, the legal system is extremely fast, reliable and efficient, thanks to distributed sensors recording continuously objective data. Deontic logic (see e.g. McNamara 2010) allows in principle to make such artificial legal reasoning. Non-violent communication (Rosenberg 2003) is used by machines to communicate empathically and efficiently with humans. Bitcoin is a distributed and decentralized digital currency which is already in use (Nakamoto 2008).

Humans are supposed to be swarming with nano-robots, which perform all kinds of measurements and enhancements, and which are connected to the internet. For pioneering work on nanomedicine, see for example (Freitas Jr 1999; 2003). In particular, mechanical artificial red cells were conceived by Freitas (1998). A dolphin speaker has recently been developed (Mishima et al. 2011). Beaming the shape of a shark might be possible, if dolphin's "sono-pictorial" language of communication would be confirmed (Kassewitz 2011). We saw earlier that surgery at distance has already been performed. An accelerated differentiation of stem cells to produce desired cells is fictional, but would be very useful, as shown in this story. Artificial viruses are likely to be used in the future given rapid and promising progress in this area (see e.g. Mastrobattista et al. 2006).

The "Association for Biological Human Rights" is fictional. But it is arguable that a too strong attachment to humans as a strictly biological (contrasted to technological) species might hinder long-term socio-technological progress.

A 3D printed bomb is also a serious threat, and the security of 3D printers should be of high concerns. The option that a coalition of artificial agents could perform such a criminal and symbolic action is a classical theme about human-machine rivalry in fiction. It also raises the following issue: could a swarm of artificial agents have a will and an agenda of its own?

This story depicted a state of high human-machine interaction and integration. But if we extrapolate the trend of externalization, fully externalized cognition doesn't need humans anymore. So, we can start to look at what intelligence would do, even if it is not grounded in human bodies.

3 - The Global Brain or the emergent whole

How will the externalized cognitive parts coordinate to form a globally distributed intelligence, or global brain? This is a very difficult question, and it presupposes that we know what we mean with "cognitive parts". A superorganism such as an ant colony or a human body consists of simple elements, ants or cells. By analogy, what will the elementary parts which compose the global brain be? Will the elementary parts be humans? Machines? Biocyborgs? A mix of the three? Or should we focus on more abstract properties, which can be embedded in different substrates? Should the elementary parts be living systems (Miller 1978), performing nineteen different functions on matter, energy and information? How intelligent should the parts be for the global brain to function? Is Artificial General Intelligence (AGI, see e.g. Ben Goertzel and Pennachin 2007) a requirement, or could the different parts be differentiated and locally relatively stupid? In this section, we propose some speculations to explore these questions.

For sure, the parts will not be humans alone, since we are already so much dependent on machines. The framework of externalizing cognition we introduced clearly suggests that we experience a human-machine symbiosis, and that it will continue for a while.

The challenge for machines is to perform better in real-world physical action. Moravec (1993) already saw the respective strengths and weaknesses of humans and machines, when he wrote that in arithmetic, "today's average computers are one million times more powerful than human beings", while "in perception and control of movement in the complex real world, and related areas of common-sense knowledge and intuitive and visual problem solving, today's average computers are a million times less capable." For practical purposes, this means that a human-machine symbiosis is suitable, since humans and machines have complementary strengths and weaknesses. As long as this dichotomy stays, the human-machine symbiosis will be the way for continuing techno-sociological progress. But how long will it last? In the context of space colonization, Hans Moravec (1979) wrote:

> The first societies in space will be composed of co-operating humans and machines, but as the capabilities of the self-improving machine component grow, the human portion will function more and more as a parasitic drag. Communities with a higher ratio of machines to people will be able to expand faster, and will become the bulk of the intelligent activity in the solar system. In the long run the sheer physical inability of humans to keep up with these rapidly evolving progeny of our minds will ensure that the ratio of people to machines approaches zero, and that a direct descendant of our culture, but not our genes, inherits the universe.

We can strengthen Moravec's argument even without taking into account the constraint of space colonization. Human bodies may become maladapted and energetically wasteful compared to ever smaller and faster information processing software and hardware. Our present and future technologies are threatening to make humans an endangered species (Joy 2000). In the near future, humans may become like dinosaurs as much as robots would become like mammals, better adapted to a new environment.

But is it really a problem? This perspective hurts our feelings and values because we have not learnt to entertain non-anthropocentric cosmic or universal values such as evolutionary, developmental or thermoethical values (for some attempts, see e.g. Ben Goertzel 2010; Vidal 2014a, chap. 10).

Even if we admit that humans may become unfit, we should not throw out the baby with the bath water and dismiss biology altogether. The possibility of a symbiosis between biology and machines constitutes the field of bionics. We could imagine a world of insect cyborgs, organizing in swarms, globally connected and able to perform any physical task. As a matter of fact, coupled machine-insect interfaces are successfully being developed (see e.g. Bozkurt et al. 2008).

The biocyborg philosophy suggests that it is easier to use and hack existing biological systems rather than to build robots from scratch. Another speculative option is the development of a cyborg bacterial intelligence or "bacteriborg". Extrapolating on the progress of synthetic biology, it actually makes most sense to try to make the most of the largest existing biomass on Earth: bacteria. Indeed, in terms of cells, the bacterial population is estimated at 5×10^{30} (Whitman, Coleman, and Wiebe 1998) while the number of cells of multicellular organisms is negligible in comparison. The idea here would be to genetically re-program bacteria, or to augment bacterial intelligence with nano (or smaller) technology, to transform the Earth into a planetary bacterial superorganism which would prove more adaptive and more intelligent than a network of humans and machines. As strange as it may seem, the incentive to choose this course of action may be purely economical. Why build from scratch (nano)robots able to intelligently communicate and collaborate on a global scale, if the same is doable by hacking readily available bacteria?

Let us now turn to a possible transition to a postbiological future. Why would that be advantageous evolutionarily? We already mentioned Moravec's argument that machines would be better adapted for space missions. In particular, they could be built to withstand very high temperatures. I mean hundreds or thousands of degrees, something impossible for life as we know it. Increasing computational speed also means increasing heat production and temperature. Unless the efficiency of cooling mechanisms finds its Moore's law, computational speed would be limited by temperature. But to understand why this point might be crucial, we need to take an even broader cosmic perspective.

4 - A Cosmic Perspective

you cannot understand one scale of system
unless you see its role in the next larger scale.

(Odum 2001, 118)

What is the next larger scale on which our planetary society is embedded? It is the cosmos. Following the insight of american ecologist Howard T. Odum, we must understand this larger cosmological context. This is motivated not by the sometimes entertaining activity of fantasizing about our far-future, but by the value of a top-down perspective, which may allow us to gain insight into where the smaller scale progresses are headed to. In Section 1, we analyzed and illustrated the logic of cognitive externalization. This was a bottom up approach. Let us now take a complementary top down cosmological approach, which promises to unveil a general evolutionary dynamic.

The overwhelming trend in technological progress is a *quest for inner space*. Indeed, our abilities to manipulate smaller and smaller entities have led and are leading to major progresses in information and communication technologies, biotechnologies and nanotechnologies (Barrow 1998; Vidal 2011; Smart 2012; Vidal 2013; Vidal 2014a). I have dubbed this trend the *Barrow scale* of civilizational development. It adds to the *Kardashev* (1964) *scale* of energy consumption. If we look at physical constraints, the quest for inner space has no reason to stop at nanotechnologies (10^{-9}m). Indeed, we could in principle develop pico- (10^{-12}m), femto- (10^{-15}m), atto- (10^{-18}m), zepto- (10^{-21}m), yocto- (10^{-24}m) technologies, and even well beyond, up to planck-technologies (10^{-35}m). The core idea with the quest for inner space is to organize matter at its smallest possible scale, and to compute with the smallest information markers allowed by physics (Bremermann 1982). This is plausible because the material support on which computers operate doesn't need to be fixed. As a matter of fact, it has already changed five times since their invention. Kurzweil (2005, chap. 3) reminds us that computers went through several hardware revolutions, from electromechanical calculators, relay-based computing, vacuum tubes, discrete transistors to today's integrated circuits. So, it seems very likely that material substrate changes will continue at a fast pace.

Additionally, material innovations account for about 2/3 of progress in computation (Magee 2012).

Quantum computer scientist Seth Lloyd (2000) explored the limits of computation and described an ultimate computer, working at a quantum scale and using reversible operations. Extrapolating Moore's law, he predicted that we would reach ultimate computation in 250 years –a very unlikely outcome for sure. First, because even such a computer would need to have error-correction mechanisms, which requires 4 x 10^{26} watts of power for a one kilogram ultimate computer. This amount of energy is more than the total output of the Sun (3.8 x 10^{26} watts)! So even if we would be able to capture all the energy of the Sun (e.g. with a Dyson (1960) sphere), it would not be enough to power just one kilogram of such an ultimate computer. Furthermore, extrapolating our global energetic consumption, we would reach the ability to harness the whole energetic output of the Sun in about 3000 years (Kardashev 1964). It means that in the best possible scenario, where we keep the pace to follow Moore's law, and where we manage to continually consume exponentially more and more energy, we would still be slowed down by the availability of free energy. Could other intelligent civilizations in the universe have been confronted to these issues, and found a solution?

In recent years, I developed a scenario for the far future of civilizations, extrapolating both the Kardashev scale and the Barrow scale (Vidal 2011). To my own surprise this led to a testable hypothesis, the *starivore hypothesis*, which claims that some known interacting binary systems are actually advanced extraterrestrial life, displaying macroscopic metabolic features (Vidal 2014a, chap. 9; 2013). In the context of our discussion, let us illustrate the putative energy use of a civilization organizing around black holes, and using them as ultimate computers. Concretely, if we take the binary microquasar SS443, it would need about 8 x 10^{57} watts if all its mass would be used to power an ultimate computer. Astrophysicists come to a close result, where 5.65 x 10^{55} watts are consummed via accretion (Fabrika, Abolmasov, and Karpov 2006, assuming perfect mass to energy conversion). Of course, this interpretation remains very speculative and difficult to prove, which is why I set up a prize to stimulate proof or disproof of the starivore hypothesis (Vidal 2014a).

5 - Conclusion

We explored the consequences of externalizing memory, computation, hearing, vision, brainstorming, reasoning, navigation, emotions and actions. Most importantly, we can foresee that new technologies will emerge as these different externalized functions are combined, and integrated with humans. Starting from known functions of our local brains, the analysis of their progressive externalization helps to better understand, foresee and facilitate the emergence of a globally distributed intelligence, best conceptualized as a global brain. With hindsight we could see the externalization and distribution of cognition as an act of evolutionary generosity, leading to the next level of intelligence on planet Earth. Indeed, we are currently externalizing, sharing and distributing our intelligence.

We critically discussed the future of human-machine symbiosis, to show the possibility of other biology-machine symbioses, such as insect-cyborgs or “bacteriborgs”. We showed that the pace of technological progress in computation is much faster than the pace of increasing energy consumption, and concluded that an energetic bottleneck is nearer than a singularity. Externalizing cognition started the transition from human-human interaction to human-machines interactions, but the cosmic perspective also suggest that intelligent life could re-organize matter to perform ultimately efficient intelligent information processing within densities, temperatures and energy scales alien to life as we know it.

6 - Questions and Answers

Q1) Ben Goertzel:

How do you react to Hugo de Garis's prediction of an "artilect war" type scenario, in which some powerful forces see the coming obsolescence of the legacy human condition as a threat, and seek to stop it from happening via force? What if some major governments are taken over by humans who are highly threatened by the various amazing developments you depict? Then, could one perhaps have an armed conflict between nations, reflecting the "Terran vs. Cosmist" dynamic that de Garis proposes? If you think this is implausible, could you articulate why?

In the fictional section 2 of my paper, I mentioned a possible conflict between the "Association for Biological Human Rights" and maybe AI agents. I want to stress that it was only for narrative purposes, as I think such conflicts are implausible, or at worst would not last long.

I do agree with much of the long-term outcomes predicted by de Garis, namely, a future where cyborgists and then cosmists become prevalent. I do also agree with de Garis's vision of hyper artilects, which fits well with the starivore hypothesis. But I fail to see any good argument showing how and why a war would play a role towards this future. Let me outline why I think such a war is implausible.

First, thinking about war presupposes a win-lose or lose-lose mental attitude. At work and in life, I focus on a win-win attitude, from which creative problem solving springs. So, de Garis's war prediction is indeed diametrically opposed to my synergetic scenario of increasing win-win interactions between humans, and artificial software or hardware agents. As the poet William Stafford famously wrote, "every war has two losers". This is not only poetry, but also a well-known result of evolutionary game-theory. The winner might benefit temporarily, but agents who make war weaken themselves at best, or simply annihilate each others. In the long run, surviving agents are the ones who reciprocate, cooperate and forgive (Axelrod 1984). So, de Garis's scenario doesn't take into account this evolutionary dynamics.

De Garis' reflections about war are also uninformed by the mechanisms leading to war and by the overwhelming empirical evidence of decreasing violence (Pinker 2011). The war predicted is naive, because our technologies are not in competition with us. Technologies do actions we can't do or we don't want to do. So, they do help us, and technology is useful. Without competition, there is nothing to fear, and no reason to start a war, from either side. To me it is similar to predicting in 1284 that the invention of wearable glasses would lead to a major war. People who don't see well refuse to become cyborgs with glasses because it would denature their true selves. However, people wearing glasses, because of their better sight, become much more efficient in acting in the world. Therefore, before they see too clearly and get all equipped with glasses, people without glasses would attempt to kill them. Seriously?

An other argument that de Garis puts forward is the IQ difference. In contrast to what de Garis argues, the IQ difference will not make a species dominance debate. First, it must be said that the IQ difference is not an absolute order. As Francis Heylighen (2011b) elaborates:

> Intelligence, like complexity, is at best represented mathematically as a partial order: for two random organisms A and B (say, a hedgehog and a magpie), A may be more intelligent than B, less intelligent, or equally intelligent, but most likely they are simply incomparable. A may be able to solve problems B cannot handle, but B can find solutions that A would not have any clue about.
>
> For such a partial order, it is impossible to develop a quantitative measure such as an IQ, because numbers are by definition fully ordered: either IQ (A) < IQ (B), IQ (A)>IQ (B), or IQ (A)=IQ (B). IQ only works in people because people are pretty similar in the type of problems they can in principle solve

Let's grant that robots develop up to the point where they become similar enough to humans so that their IQ may be compared. It will be too obvious who is smarter and who will win the debate, like it is obvious that we are smarter than apes. We don’t argue nor make war with apes. We don’t ask them if they agree that we destroy or protect the forest they live in. We just do it. However, it raises the questions of what to do with less dominant species, as with the issues of animal rights. Should we make them dissapear? Should we leave them at peace? Should we uplift them to higher levels of intelligence? These are ethical questions, whose alternatives are rich and complex. So, even if there is a superiority, it would not even count as war, so much the difference in cognitive and physical capacity would be different.

I would like also to pinpoint that the distinction between “Terrans” and “Cosmists” is naive. At first sight, it seems like a plausible distinction, but it doesn’t resist closer analysis, because the boundary is fuzzy. De Garis seems to put the cyborgists on the side of cosmists. But the matter is much more complicated. Let’s say I’m a radical Terran. At which point would I want kill my grandmother? If she wears glasses to be able to read? If she uses a phone? A smartphone? A pacemaker? Of course, it is easy to see that all the technologies she uses are for her own benefit and do not threaten her. So, it would really make no sense to kill my beloved and augmented grandma. We don’t make war against our grandmas who choose not to use smartphones. At worst, we just interact less with them than with our hyperconnected friends. This example shows that there would be no clear opponents, so no clear reason to start a war.

Even granting all the arguments, predicting a war is the worst thing a futurist can do. Contrary to what Prof. de Garis holds, being pessimist or optimist is extremely relevant for futurology. We build the future, we don’t just predict it like an apple falling from a tree. Futurists usually [sic!] know the power of self-fulfilling prophecies, and the impact of their models on reality. A self-fulfilling prophecy says that when authorities say something will happen, it happens simply because many people believe it. I think it is the professional duty of futurists to remain optimists. Of course, it doesn’t mean we have to be naive optimists, and we still should anticipate and avoid dangers, but these are not the ones we should emphasize if we want to build a better future.

I don’t exclude a priori the eventuality of a future war. If there would be a war, I would rather bet on a more traditionally motivated war, where there is a clear fight for resources, or any other geopolitical motivations.

Despite all the counter-arguments above, let’s assume de Garis is right, and a war is indeed coming. Shouldn't he do whatever he can to avoid it? How can we make it a self-defeating prophecy?

Our short analysis shows that everything goes against de Garis' war scenario. Basic knowledge of evolutionary dynamics and game theory shows that complex systems tend towards more synergetic or win-win interactions, and not win-lose. The empirical evidence strongly supports a decrease of violence, not an increase. The mental attitude behind war is non-creative and non-constructive, and since a futurist also influences the unfolding of the future, it is his duty to emphasize positive futures, and to show how to avoid negative ones.

De Garis' worst case scenario is possible at most, but arguably highly improbable and obviously highly undesirable. There are all reasons to consider de Garis' scenario as material for naive science fiction, but not for professional futurology. To sum up, de Garis' vision can be taken as seriously as blockbuster movies entitled "Smartphones versus Humans" or "Grandmas versus Terminators".

Q2) Ted Goertzel:

You suggest that, with advances in technology, humans may become maladapted like dinosaurs that were unable to adapt to compete with mammals. But I understand that the dinosaurs were able to contain the threat from mammals and remain dominant until they were destroyed by a natural disaster, either a comet or asteroid or massive volcanism. Might not humans similarly suppress or control technologies that threaten their dominance on the planet?

The analogy is interesting, because it's not the mammals that kill the dinosaurs. It's simply that mammals were able to better cope with a catastrophic event -asteroid impact and subsequent volcanism. Dinosaurs never intentionally controlled or contained a mammal threat.

I think it is hard to suppress or control technologies, even if they are threatening us. Let us consider the greatest technological backbone, the physical internet network. Is it only possible to suppress it today? It would be very hard, arguably as hard as suppressing electrical networks or hydraulic networks. More fundamentally, it would be a regression.

Technological progress is an arms race. The countries slowing down or suppressing technological progress for ethical or political reasons will get outcompeted by countries without such restrictions. So, suppressing technologies is certainly not a good strategic move, and the history of technology strongly supports this view. Additionally, it will certainly remain hard if not impossible to reach agreement on the suppression or control of technologies on a global scale.

I would say the smartest way is to accompany technological progress as it appears, with debates and reflections, without slowing it down. You could imagine setting up global ethical institutes, and for some applications, it might be a good thing. But the mission would not be easy to define. What do we want to protect? Should we restrict genetic engineering and cyborgisation in name of preserving the purity of the human species?

Who should become the dominant species? Humanity at all price or the smartest and most adaptive creatures? As we saw, even if we reply with "humanity", it's not clear where we set its boundaries. This is why I think broader evolutionary, developmental and thermodynamic values and ethics are much needed (Vidal 2014a, chap. 10). They open ways to reason ethically without the fallacy of putting humanity at the center.

Q3) ***Ben Goertzel:***

In the future you project, the nature of intelligence will gradually shift from its current individual mind/body focused state, to a situation where intelligence is more associated with global distributed networks…. What do you think this will do to the experience of being a mind? Modern psychology is highly focused on the ego and the self, on individual ambitions and goals and possessions and desires and dreams. Do you think this will change? If so, what will replace it? What will it be like to be a distributed network mind – or a part of such a highly richly interconnected "mindplex"-- rather than an ego-centered individual?

There are two different questions you are asking.

1) What is it like *to be part of* the global brain?
2) What is it like *to be* the global brain?

Regarding the first question, I am not a psychologist, but I am sure that the effects of our networked society on our psychology are disruptive. I can only skim over the question in this reply. Nicholas Carr (2011) has explored what the internet does to our brain, although emphasizing more negative side-effects. Psychologists study internet addiction (e.g. Young 1998), which means to me that we need more and more to be always connected. We do already experience our mind as hyperconnected to the internet.

We need to process more and more information every year if not every month or week. This need to deal with more information can be stressful and distracting: too many videos to watch, too many things to read, too many emails to answer to. To cope with information overload, we need to be able to process it quickly and efficiently. Fortunately, there is no obligation to use our biological brains only. If you are a boss, you can distribute or share the workload with peers (humans). But for the rest of us, we can delegate simple tasks to human agents worldwide, by distributing microwork in crowsourcing internet marketplaces. Maybe the visually impaired persons who use daily microworkers to make sense of their environment consider such workers really as part of their self?

But we can also use software agents, such as spam email filters. We must augment our processing speed capacity, and we must improve our information filters. Such algorithms are likely to become more and more essential to deal with everyday life, to a point where everybody will have a personal operating system, informed by big data, knowing our values and preferences, and guiding us. This could be implemented as a conversation, like in Spike Jonze's movie "Her". However, the conversational interface is optional and it could simply become a smarter "I", which selects the relevant information.

In developmental psychology, the definition of the self and ego changes as the individual grows from childhood to adulthood. As we continue to grow in our networked society, we care increasingly for global issues, as if they would impact our own local lives. Look at the number of non-governmental organizations. People do care about planet Earth, it's not anymore a New Age trend. We are broadening our circles of compassions to the Earth as a whole, and I find this the most beautiful development of globalization.

As a side note, I see no reason to stop at the boundaries of our tiny planet. I have extrapolated this trend of increasing compassion (Vidal 2014a, sec. 10.3.3; Vidal 2014b), where beings would gradually care for the Earth, the Sun, the Galaxy, the Universe and finally the process of evolution itself.

To speculate about the ego, ambitions, goals, possessions, desires and dreams of the global brain is a stimulating speculative exercise, which would require a textbook entitled "Global Brain Psychology". I attempt in my answer to Ted's question below to account for the global brain's cognitive development below. Let me first say a word about the scientific study of the global brain's subjective experience.

To describe the inner experience of the global brain is nearly impossible, as it is to describe the experience of a tree. Actually, even guessing the inner feelings of other humans, we are often wrong. So, there is no doubt that guessing the feelings of the global brain is orders of magnitude harder.

Yet, even if we give up the idea of trying to "feel the Earth", we can still study it, as we can study trees or human cognition. The fact that we can not easily project our subjective experience doesn't forbid objective scientific work, it just makes it less intuitive. A typical concrete example is to measure the global mood through Twitter, which can be used, amongst other, to predict the stock market (Bollen, Mao, and Zeng 2011).

Q4) Ted Goertzel:

You do not offer a prediction of how soon a globally distributed intelligence, or global brain, will occur. Of course, a degree of globally distributed intelligence already exists, but its capabilities are less, or at least quite different, from those of the human brain. Can you offer some predictions as to how long it will take for a global brain to emerge? Or, even better, how long it will take for key steps towards a full global brain to emerge?

I see the global brain more as an ongoing transition rather than a precise event. As you mention, the transition has started for sure. I am not an empirical futurist, so I apologize that I will not attempt to go into the business of precise qualitative predictions. Qualitative outcomes depend on many variables such as political choices, priorities, resource allocations, or markets. So they may be wrong by few decades. To illustrate this, I would be open to the possibility that the science-fiction story I proposed in section 2, "A day in 2060", could happen in 2030 or 2080. But different quantitative estimates of the singularity do converge, and this is remarkable, because while the estimates start from different assumptions, they converge to a similar critical date (Magee and Devezas 2011). To be on the safe side, I would say that a major transition will have occurred before the end of the century.

That said, I can offer a few reflections on key steps towards a global brain, out of which *qualitative* predictions might be made.

As you note, the intelligence capabilities of the embryonic global brain are different from those of the human brain. This is an important point, because it shows that different kinds of intelligences can co-exist at different scales, and benefit from each others. However, the analogy of the global brain and evolutionary psychology is a useful framework for thinking and maybe predicting the future of the globally distributed intelligence (Vidal, Heylighen, and Weinbaum 2013).

Although it is an approximation and a simplification, we can think of the global brain's cognitive development in three phases: reptilian, limbical and neocortical.

The reptilian brain reacts and controls primitive bodily functions such as hunger, fear, fight or flight. It connects the nervous system to limbs and triggers instinctual behaviors. At the scale of the global brain, this could correspond to a *planetary nervous system* (Helbing et al. 2012). The main goal here is to deal with global risks, such as tsunamis, epidemics, or volcanic eruptions, thanks to distributed sensors over the globe, connected to effectors, which can be humans or machines.

The limbic system is involved in long-term memory, emotion and motivation. At a global scale, these functions are performed by social networks, which weave a rather emotionally-driven network, and can lead to spontaneous eruptions, such as the Arab Spring revolutions. Wikipedia would also be part of this limbic development, acting as a dynamic and collective memory.

The neocortex is evolutionarily the most recent part of the brain, essential to human language, abstract thinking, imagination or awareness. At the global brain scale, it corresponds to two challenges: building global simulations, and fostering collective reasoning and decision-making. With the help of computer simulations fed with big data, we start to better understand how to tackle global and complex issues. We thus make wiser decisions, and this constitutes a kind of global imagination or dreaming faculty. The process is indeed similar to the faculty of our minds to play possible scenarios, and to dreaming, when it helps to consolidate the most important memories. I mentioned in my paper the importance of collective reasoning through its externalization and distribution. If we use data both from distributed sensors and computer simulations, and add effective methods for collective reasoning and decision making, it is certain that smarter decisions will be made on a global scale.

These three levels constitute a roadmap for the development of the global brain, which maps remarkably well with the three components of the FuturICT project; respectively the *planetary nervous system*, the *living earth simulator* and the *global participatory platform* (Helbing et al. 2012). In my eyes, we are still working on the reptilian and limbic parts. It is an interesting question whether and in how far we can or should develop the three systems sequentially or in parallel.

But the whole is more than the sum of the parts, so we can ask: what would the global brain then want to do, in order to foster its cognitive development? The answer is simply: the same as us, become part of a social group, interact, learn from others. So, as strange as it may seem, I predict that as the global brain unfolds, it will become more and more pressing to find other global brains in the universe to communicate with and to learn from. To develop the "self" of an intelligent planet, you need to find others with whom you can communicate and with whom you can compare yourself to.

As the co-founder of the Evo Devo Universe community (Smart 2008; Vidal 2010), I believe that in the midst of unpredictable and chaotic evolution, larger developmental trends can be identified in our universe. So I would like to add a few thoughts on this cosmic perspective, which is an important addition next to the five approaches to predicting long-term trends you identified (Ben Goertzel and Goertzel 2014, this volume).

One aspect of this cosmic perspective is precisely astrobiology, or the search for life elsewhere in the universe. As Kardashev (1978) noted, finding extraterrestrial intelligence (ETI) might give us a glimpse of our possible future on astronomical timescales. We would have other examples, that we can choose to follow … or not. Even if we are alone in the universe, or if we don't find anything, the exercise still has immense value because we must project our civilization into the future to search for other advanced ETIs.

In light of the cosmic perspective, there is another clear prediction. *The global brain will be hungry*. This is a prediction of fundamental importance that futurists seem to overlook. While human population is stabilizing, there is a population explosion of transistors which will need energy. Information communications technologies (ICTs) use about 10% of the world electricity generation (1500 TWh, see Mills 2013). If the global brain is like the human brain, then the consumption could raise up to 20%, which is the ratio of energy used by the human brain (Raichle and Gusnard 2002).

In a cosmic setting, the global brain is just the beginning. Regarding what could happen after the global brain, I invite you to consider the starivore hypothesis, and the hypothetical possibilities of major transitions towards global organizations on nuclear scales (white dwarfs), sub-nuclear (neutron stars), and space-time structure (black holes). These small-scales and high-density states of matter turned into intelligent organizations would necessitate increasingly more and more energy, nothing less than the total energy of a star.

Q5) Ben Goertzel:

I love your speculations about starivores and so forth. I myself have spent a bit of time musing about how one might construct machinery (of a sort) and intelligent systems within degenerate matter such as quark stars or quark-gluon plasmas – going beyond nanotech to femtotech, in this way. I note that you say you've "set up a prize to stimulate proof or disproof of the starivore hypothesis"…. While somewhat peripheral to the main focus of this book, this is sufficiently intriguing to me that I can't pass up asking about it. What kind of proof or disproof do you think might be found in the near future? How would you envision concretely demonstrating the existence of intelligence in far-off star systems? Via some sort of conventional SETI-like analysis of astronomical data from these systems, I'd suppose? Also, I wonder if you have any thoughts about the internal nature of these starivore minds? What is it like to be a starivore?

I'm glad that you have also considered nuclear or sub-nuclear substrates as a basis for intelligent organization.

Proving the existence of extraterrestrial intelligence will not be easy and clear-cut, whether it is the starivore hypothesis or another idea. The question we need to address is: what kind of proof would convince the most skeptical scientists? There are only three such kinds of proofs: prediction, prediction and prediction. I mean, predicting matter, energy and information patterns.

Concretely, this could mean predicting movement patterns, e.g. showing that the direction of some moving white dwarfs is not random, but targeted towards the nearest food source, that is, a star. Predicting energy patterns, we can assume that, as living systems, starivores should follow one of the most universal biological laws, Kleiber's law (DeLong et al. 2010). Predicting information patterns, X-ray and millisecond pulsars are particularly intriguing and promising. I refer the curious reader to Chap. 9 of my (2014a) book for more details on possible predictions.

What I have in mind with these projects leading to predictions is simply the regular hypothetico-deductive way of science. The logic goes as follows. Let's assume that such systems are artificial. Can we make new predictions that we can't derive when we assume that the systems are natural? If we make such predictions, such as the ones above, and if they become verified, then we build the case that such binary stars are extraterrestrials. If not, it will tend to falsify the starivore hypothesis.

Conventional SETI only matches the third category, the prediction of information patterns. Even there, we can distinguish at least two distinct issues:

- proving that a signal is intelligent via some abstract complexity measure such as Kolmogorov (e.g. Li and Vitányi 1997) or logical depth (Bennett 1988)
- decoding an extraterrestrial message.

I believe this second challenge is way harder than the first. To sum up, I envision and hope that 5-10 successful and different predictions like these would make skeptics scratch their heads, and trigger the scientific community to look at the starivore hypothesis more and more seriously and intensely.

What's inside a starivore mind? The best way to answer this question would be to ask them. But it would take a few decades or centuries to make a Q&A with them. So, in the meantime, let us speculate. As a minimal hypothesis, we can assume that they are living organisms. As such, they first want to ensure basic survival needs, first and foremost, eat stars. Of course, eating a star might be a science and an art, in addition to being an engineering feat. So how to eat a star might be key questions in starivore's minds. When this basic energetic need is fulfilled, they would use their computing and cognitive power to explore and understand their surroundings. But, to understand a complex system, there is often no other shortcut than to run a simulation of it. So, they would run very advanced versions of the "living earth simulator", simulating the whole universe and not only their own local world.

I imagine that they would have went over their global brain transition and act as a whole and coherent living system, beyond local frictions and dissensions. In the long run, their self or ego would even transcend their own star-eating existence, realizing the futility of eating stars all days, if the universe is doomed to dissipate all its energy, as predicted by the second law of thermodynamics. So, they would turn into the noble mission of securing the life of the cosmos, maybe first through stellar and galactic rejuvenation, and then through the making of a new universe (Vidal 2014b).

7 - Acknowledgments

I thank Viktoras Veitas and David Weinbaum for helpful critical feedback, as well as Ted and Ben Goertzel for their challenging questions.

8 - References

Anderson, Michael L. 2003. "Embodied Cognition: A Field Guide." *Artificial Intelligence* 149 (1): 91–130. doi:10.1016/S0004-3702(03)00054-7.

Anvari, Mehran, Craig McKinley, and Harvey Stein. 2005. "Establishment of the World's First Telerobotic Remote Surgical Service." *Annals of Surgery* 241 (3): 460–64. doi:10.1097/01.sla.0000154456.69815.ee.

Axelrod, Robert M. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Baldwin, Peter, and David Price. 2008. "Debategraph." http://debategraph.org/.

Barrow, J. D. 1998. *Impossibility: The Limits of Science and the Science of Limits*. Oxford University Press, USA.

Bennett, C.H. 1988. "Logical Depth and Physical Complexity." In *The Universal Turing Machine: A Half-Century Survey*, edited by R. Herken, 227–57. Oxford University Press. http://www.research.ibm.com/people/b/bennetc/UTMX.pdf.

Berners-Lee, T. 1999. *Weaving the Web*. HarperBusiness New York.

Bigham, J. P, C. Jayant, H. Ji, G. Little, A. Miller, R. C Miller, R. Miller, et al. 2010. "VizWiz: Nearly Real-Time Answers to Visual Questions." In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*, 333–42.

Bollen, Johan, Huina Mao, and Xiaojun Zeng. 2011. "Twitter Mood Predicts the Stock Market." *Journal of Computational Science* 2 (1): 1–8. doi:10.1016/j.jocs.2010.12.007.

Bozkurt, A., R. Gilmour, D. Stern, and A. Lal. 2008. "MEMS Based Bioelectronic Neuromuscular Interfaces for Insect Cyborg Flight Control." In *IEEE 21st International Conference on Micro Electro Mechanical Systems, 2008. MEMS 2008*, 160–63. doi:10.1109/MEMSYS.2008.4443617.

Bremermann, Hans J. 1982. "Minimum Energy Requirements of Information Transfer and Computing." *International Journal of Theoretical Physics* 21: 203–17. doi:10.1007/BF01857726.

Carr, Nicholas G. 2011. *The Shallows: What the Internet Is Doing to Our Brains*. New York: W.W. Norton.

Chorley, Martin. 2013. "Augmented Conversation - Summer Project 2011." *Martin Chorley*. Accessed July 9. http://martinjc.com/research/student-projects/augmented-conversation/.

Clark, Andy. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

Cohen, Daniel L., Jeffrey I. Lipton, Meredith Cutler, Deborah Coulter, Anthony Vesco, and Hod Lipson. 2009. "Hydrocolloid Printing: A Novel Platform for Customized Food Production." In *Solid Freeform Fabrication Symposium*. http://edge.rit.edu/edge/P10551/public/SFF/SFF%202009%20Proceedings/2009%20SFF%20Papers/2009-71-Cohen.pdf.

DeLong, John P., Jordan G. Okie, Melanie E. Moses, Richard M. Sibly, and James H. Brown. 2010. "Shifts in Metabolic Scaling, Production, and Efficiency Across Major Evolutionary Transitions of Life."

Proceedings of the National Academy of Sciences 107 (29): 12941 – 12945. doi:10.1073/pnas.1007783107.
Dror, Itiel E., and Stevan Harnad. 2008. "Offloading Cognition onto Cognitive Technology." In *Cognition Distributed: How Cognitive Technology Extends Our Minds*, 1–23. http://arxiv.org/abs/0808.3569.
Dyson, F. J. 1960. "Search for Artificial Stellar Sources of Infrared Radiation." *Science* 131 (3414): 1667 –1668. doi:10.1126/science.131.3414.1667.
Fabrika, S. N., P. K. Abolmasov, and S. Karpov. 2006. "The Supercritical Accretion Disk in SS 433 and Ultraluminous X-Ray Sources." *Proceedings of the International Astronomical Union* 2 (S238): 225–28.
Fischer, Gerhard. 2006. "Distributed Intelligence: Extending the Power of the Unaided, Individual Human Mind." In *In Augusto Celentano (Ed.), Proceedings of the Advanced Visual Interfaces (AVI) Conference (pp. 7–14).New*, 23–26. ACM Press. http://l3d.cs.colorado.edu/~Gerhard/papers/avi-2006.pdf.
Freitas Jr, R. A. 1998. "Exploratory Design in Medical Nanotechnology: A Mechanical Artificial Red Cell." *Artificial Cells, Blood Substitutes, and Immobilization Biotechnology* 26 (4): 411–30.
———. 1999. *Nanomedicine, Volume I: Basic Capabilities*. Georgetown, TX: Landes Bioscience. http://www.nanomedicine.com/NMI.htm.
———. 2003. *Nanomedecine Volume IIA: Biocompatibility.* Georgetown, TX: Landes Bioscience. http://www.nanomedicine.com/NMIIA.htm.
Goertzel, Ben. 2010. *A Cosmist Manifesto: Practical Philosophy for the Posthuman Age*. Humanity+. http://goertzel.org/CosmistManifesto_July2010.pdf.
Goertzel, Ben, and Ted Goertzel. 2014. "Introduction: Where Are We Going? How Will We Get There?" In *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*, edited by Ben Goertzel and Ted Goertzel.
Goertzel, Ben, and Cassio Pennachin, eds. 2007. *Artificial General Intelligence*. Cognitive Technologies. Berlin ; New York: Springer.
Golan, Lior, Daniella Yeheskely-Hayon, Limor Minai, Eldad J Dann, and Dvir Yelin. 2012. "Noninvasive Imaging of Flowing Blood Cells Using Label-Free Spectrally Encoded Flow Cytometry." *Biomedical Optics Express* 3 (6): 1455–64. doi:10.1364/BOE.3.001455.
Goldratt, E. M., and J. Cox. 1984. *The Goal: A Process of Ongoing Improvement*. 3rd ed. Great Barrington, MA: North River Press.
Helbing, D., S. Bishop, R. Conte, P. Lukowicz, and J. B. McCarthy. 2012. "FuturICT: Participatory Computing to Understand and Manage Our Complex World in a More Sustainable and Resilient Way." *The European Physical Journal Special Topics* 214 (1): 11–39. doi:10.1140/epjst/e2012-01686-y.
Heylighen, F. 2007a. "The Global Superorganism: An Evolutionary-Cybernetic Model of the Emerging Network Society." *Social Evolution & History* 6 (1): 58–119. doi:10.1.1.43.3443.
———. 2007b. "Accelerating Socio-Technological Evolution: From Ephemeralization and Stigmergy to the Global Brain." In *Globalization as an Evolutionary Process: Modeling Global Change*, edited by George Modelski, Tessaleno Devezas, and William Thompson,

Routledge, 286–335. London. http://pespmc1.vub.ac.be/Papers/AcceleratingEvolution.pdf.
———. 2011a. “Conceptions of a Global Brain: An Historical Review.” In *Evolution: Cosmic, Biological, and Social, Eds. Grinin, L. E., Carneiro, R. L., Korotayev A. V., Spier F.*, 274 – 289. Uchitel Publishing. http://pcp.vub.ac.be/papers/GBconceptions.pdf.
———. 2011b. Francis Heylighen on the Emerging Global Brain. Interview by B. Goertzel. H+ Magazine. http://hplusmagazine.com/2011/03/16/francis-heylighen-on-the-emerging-global-brain/.
———. 2013. *Distributed Intelligence Technologies: A Survey of Future Applications and Implications of the Global Brain*. Global Brain Institute, Working Paper.
Heylighen, F., and C. Vidal. 2008. “Getting Things Done: The Science behind Stress-Free Productivity.” *Long Range Planning* 41 (6): 585–605. doi:10.1016/j.lrp.2008.09.004. http://cogprints.org/5904/
Iandoli, Luca, Mark Klein, and Giuseppe Zollo. 2007. *Can We Exploit Collective Intelligence for Collaborative Deliberation? The Case of the Climate Change Collaboratorium*. MIT Sloan Research Paper No. 4675-08. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1084069.
Jacobs, Goff, Amy Aeron-Thomas, and Angela Astrop. 2000. *Estimating Global Road Fatalities*. Transport Research Laboratory Crowthorne. http://ministryofsafety.files.wordpress.com/2011/01/estimating-global-road-fatalities-full-report1.pdf.
Joy, Bill. 2000. “Why the Future Doesn’t Need Us.” In *Nanoethics–the Ethical and Social Implicatons of Nanotechnology*, 17–39. http://www.aaas.org/spp/rd/ch3.
Kardashev, N. S. 1964. “Transmission of Information by Extraterrestrial Civilizations.” *Soviet Astronomy* 8 (2): 217–20. http://adsabs.harvard.edu/abs/1964SvA.....8..217K.
———. 1978. “On Strategy in the Search for Extraterrestrial Civilizations.” *Russian Social Science Review* 19 (4): 27–47. doi:10.2753/RSS1061-1428190427.
Kassewitz, Jack. 2011. “We Are Not Alone: The Discovery of Dolphin Language.” http://www.speakdolphin.com/ResearchItems.cfm?ID=20.
Kittur, Aniket, Bongwon Suh, Bryan A. Pendleton, and Ed H. Chi. 2007. “He Says, She Says: Conflict and Coordination in Wikipedia.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 453–62. CHI ’07. New York, NY, USA: ACM. doi:10.1145/1240624.1240698.
Kurzweil, R. 2005. *The Singularity Is Near: When Humans Transcend Biology.* Penguin Books.
Li, Ming, and P. M. B. Vitányi. 1997. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer.
Lingley, A R, M Ali, Y Liao, R Mirjalili, M Klonner, M Sopanen, S Suihkonen, et al. 2011. “A Single-Pixel Wireless Contact Lens Display.” *Journal of Micromechanics and Microengineering* 21 (December): 125014. doi:10.1088/0960-1317/21/12/125014.
Lipson, Hod, and Melba Kurman. 2013. *Fabricated: The New World of 3D Printing*. 1st ed. Wiley.

Lloyd, S. 2000. “Ultimate Physical Limits to Computation.” *Nature* 406: 1047–54. http://www.hep.princeton.edu/~mcdonald/examples/QM/lloyd_nature_406_1047_00.pdf.

Magee, Christopher L. 2012. “Towards Quantification of the Role of Materials Innovation in Overall Technological Development.” *Complexity* 18 (1): 10–25. doi:10.1002/cplx.20309.

Magee, Christopher L., and Tessaleno C. Devezas. 2011. “How Many Singularities Are near and How Will They Disrupt Human History?” *Technological Forecasting and Social Change* 78 (8): 1365–78. doi:10.1016/j.techfore.2011.07.013.

Mastrobattista, Enrico, Marieke A. E. M. van der Aa, Wim E. Hennink, and Daan J. A. Crommelin. 2006. “Artificial Viruses: A Nanotechnological Approach to Gene Delivery.” *Nature Reviews Drug Discovery* 5 (2): 115–21. doi:10.1038/nrd1960.

McNamara, Paul. 2010. “Deontic Logic.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2010. http://plato.stanford.edu/archives/fall2010/entries/logic-deontic/.

Miller, J. G. 1978. *Living Systems*. McGraw-Hill New York.

Mills, M. P. 2013. *The Cloud Begins with Coal: Big Data, Big Networks, Big Infrastructure, and Big Power.* National Mining Assoication & American Coalition for clean Coal Electricity. http://write.americaspower.org/sites/default/files/The%20Cloud%20Begins%20with%20Coal%20Full%20Report.pdf.

Minder, Patrick, and Abraham Bernstein. 2012. “CrowdLang: A Programming Language for the Systematic Exploration of Human Computation Systems.” In *Social Informatics*, edited by Karl Aberer, Andreas Flache, Wander Jager, Ling Liu, Jie Tang, and Christophe Guéret, 124–37. Lecture Notes in Computer Science 7710. Springer Berlin Heidelberg. http://link.springer.com/chapter/10.1007/978-3-642-35386-4_10.

Mironov, Vladimir, Thomas Boland, Thomas Trusk, Gabor Forgacs, and Roger R. Markwald. 2003. “Organ Printing: Computer-Aided Jet-Based 3D Tissue Engineering.” *Trends in Biotechnology* 21 (4): 157–61. doi:10.1016/S0167-7799(03)00033-7.

Mishima, Yuka, Keiichi Uchida, Kazuo Amakasu, Yoshinori Miyamoto, and Toyoki Sasakura. 2011. “Development of Dolphin-Speaker.” *The Journal of the Acoustical Society of America* 130 (4): 2358. doi:10.1121/1.3654450.

Moravec, Hans. 1979. “Today’s Computers, Intelligent Machines and Our Future.” *Analog* 99 (2): 59–84. http://www.frc.ri.cmu.edu/~hpm/project.archive/general.articles/1978/analog.1978.html.

———. 1993. “The Universal Robot.” In , 35–41. http://adsabs.harvard.edu/abs/1993vise.nasa...35M.

Morris, Robert R, and Rosalind Picard. 2012. “Crowdsourcing Collective Emotional Intelligence.” In Cambridge, MA. http://arxiv.org/abs/1204.3481.

Nagy, Béla, J. Doyne Farmer, Jessika E. Trancik, and John Paul Gonzales. 2011. “Superexponential Long-Term Trends in Information

Technology.” *Technological Forecasting and Social Change* 78 (8): 1356–64. doi:10.1016/j.techfore.2011.07.006.
Nakamoto, Satoshi. 2008. *Bitcoin: A Peer-to-Peer Electronic Cash System*. http://bitcoin.org/bitcoin.pdf.
Odum, H. T. 2001. Interview of Dr. Howard T. Odum Interview by Cynthia Barnett.
Picard, Rosalind W. 1997. *Affective Computing*. Cambridge, Mass: MIT Press.
Pinker, Steven. 2011. *The Better Angels of Our Nature: Why Violence Has Declined*. Viking Adult.
Raichle, Marcus E., and Debra A. Gusnard. 2002. “Appraising the Brain’s Energy Budget.” *Proceedings of the National Academy of Sciences* 99 (16): 10237–39. doi:10.1073/pnas.172399499.
Robertson, Douglas S. 1998. *The New Renaissance Computers and the next Level of Civilization*. New York: Oxford University Press.
Rosenberg, Marshall B. 2003. *Nonviolent Communication: A Language of Life*. Second edition. Puddledancer Press.
Smart, J. 2008. “Evo Devo Universe? A Framework for Speculations on Cosmic Culture.” In *Cosmos and Culture*, edited by S. J. Dick. To appear. http://accelerating.org/downloads/SmartEvoDevoUniv2008.pdf.
———. 2012. “The Transcension Hypothesis: Sufficiently Advanced Civilizations Invariably Leave Our Universe, and Implications for Meti and Seti.” *Acta Astronautica*, no. 0. doi:10.1016/j.actaastro.2011.11.006.
Vidal, C. 2010. “Introduction to the Special Issue on the Evolution and Development of the Universe.” *Foundations of Science* 15 (2): 95–99. doi:10.1007/s10699-010-9176-9.
———. 2011. “Black Holes: Attractors for Intelligence?” In Buckinghamshire, Kavli Royal Society International Centre. http://arxiv.org/abs/1104.4362.
———. 2013. “Starivore Extraterrestrials? Interacting Binary Stars as Macroscopic Metabolic Systems.” *Working Paper*. http://student.vub.ac.be/~clvidal/writings/Vidal-Starivore-Binary.pdf.
———. 2014a. *The Beginning and the End: The Meaning of Life in a Cosmological Perspective*. 2014th ed. New York: Springer http://arxiv.org/abs/1301.1648.
———. 2014b. “Cosmological Immortality: How to Eliminate Aging on a Universal Scale.” *Current Aging Science*. doi:10.2174/1874609807666140521111107. http://student.vub.ac.be/~clvidal/writings/Vidal-Cosmological-Immortality.pdf
Vidal, C., F. Heylighen, and D. Weinbaum. 2013. “Evolutionary Psychology of the Global Brain and FuturICT”. European Conference on Complex Systems presented at the Satellite Meeting “Global Computing for our Complex Hyper-connected World,” Barcelona, September 19. http://prezi.com/pv9t6-ndmyy6/eccs-2013-evolutionary-psychology-of-the-global-brain-and-futurict/.
Whitman, William B., David C. Coleman, and William J. Wiebe. 1998. “Prokaryotes: The Unseen Majority.” *Proceedings of the National Academy of Sciences* 95 (12): 6578–83. http://www.pnas.org/content/95/12/6578.

Wilson, Margaret. 2002. "Six Views of Embodied Cognition." *Psychonomic Bulletin & Review* 9 (4): 625–36. doi:10.3758/BF03196322.
Young, Kimberly S. 1998. "Internet Addiction: The Emergence of a New Clinical Disorder." *CyberPsychology & Behavior* 1 (3): 237–44. doi:10.1089/cpb.1998.1.237.

9 - Argumentative maps

The following two pages are argumentative maps summarizing the main points of the paper. It uses the methodology of the theory of constraints (TOC). The first map is the problem (or "current reality tree" in TOC), the second the proposed solution (or "future reality tree" in TOC).

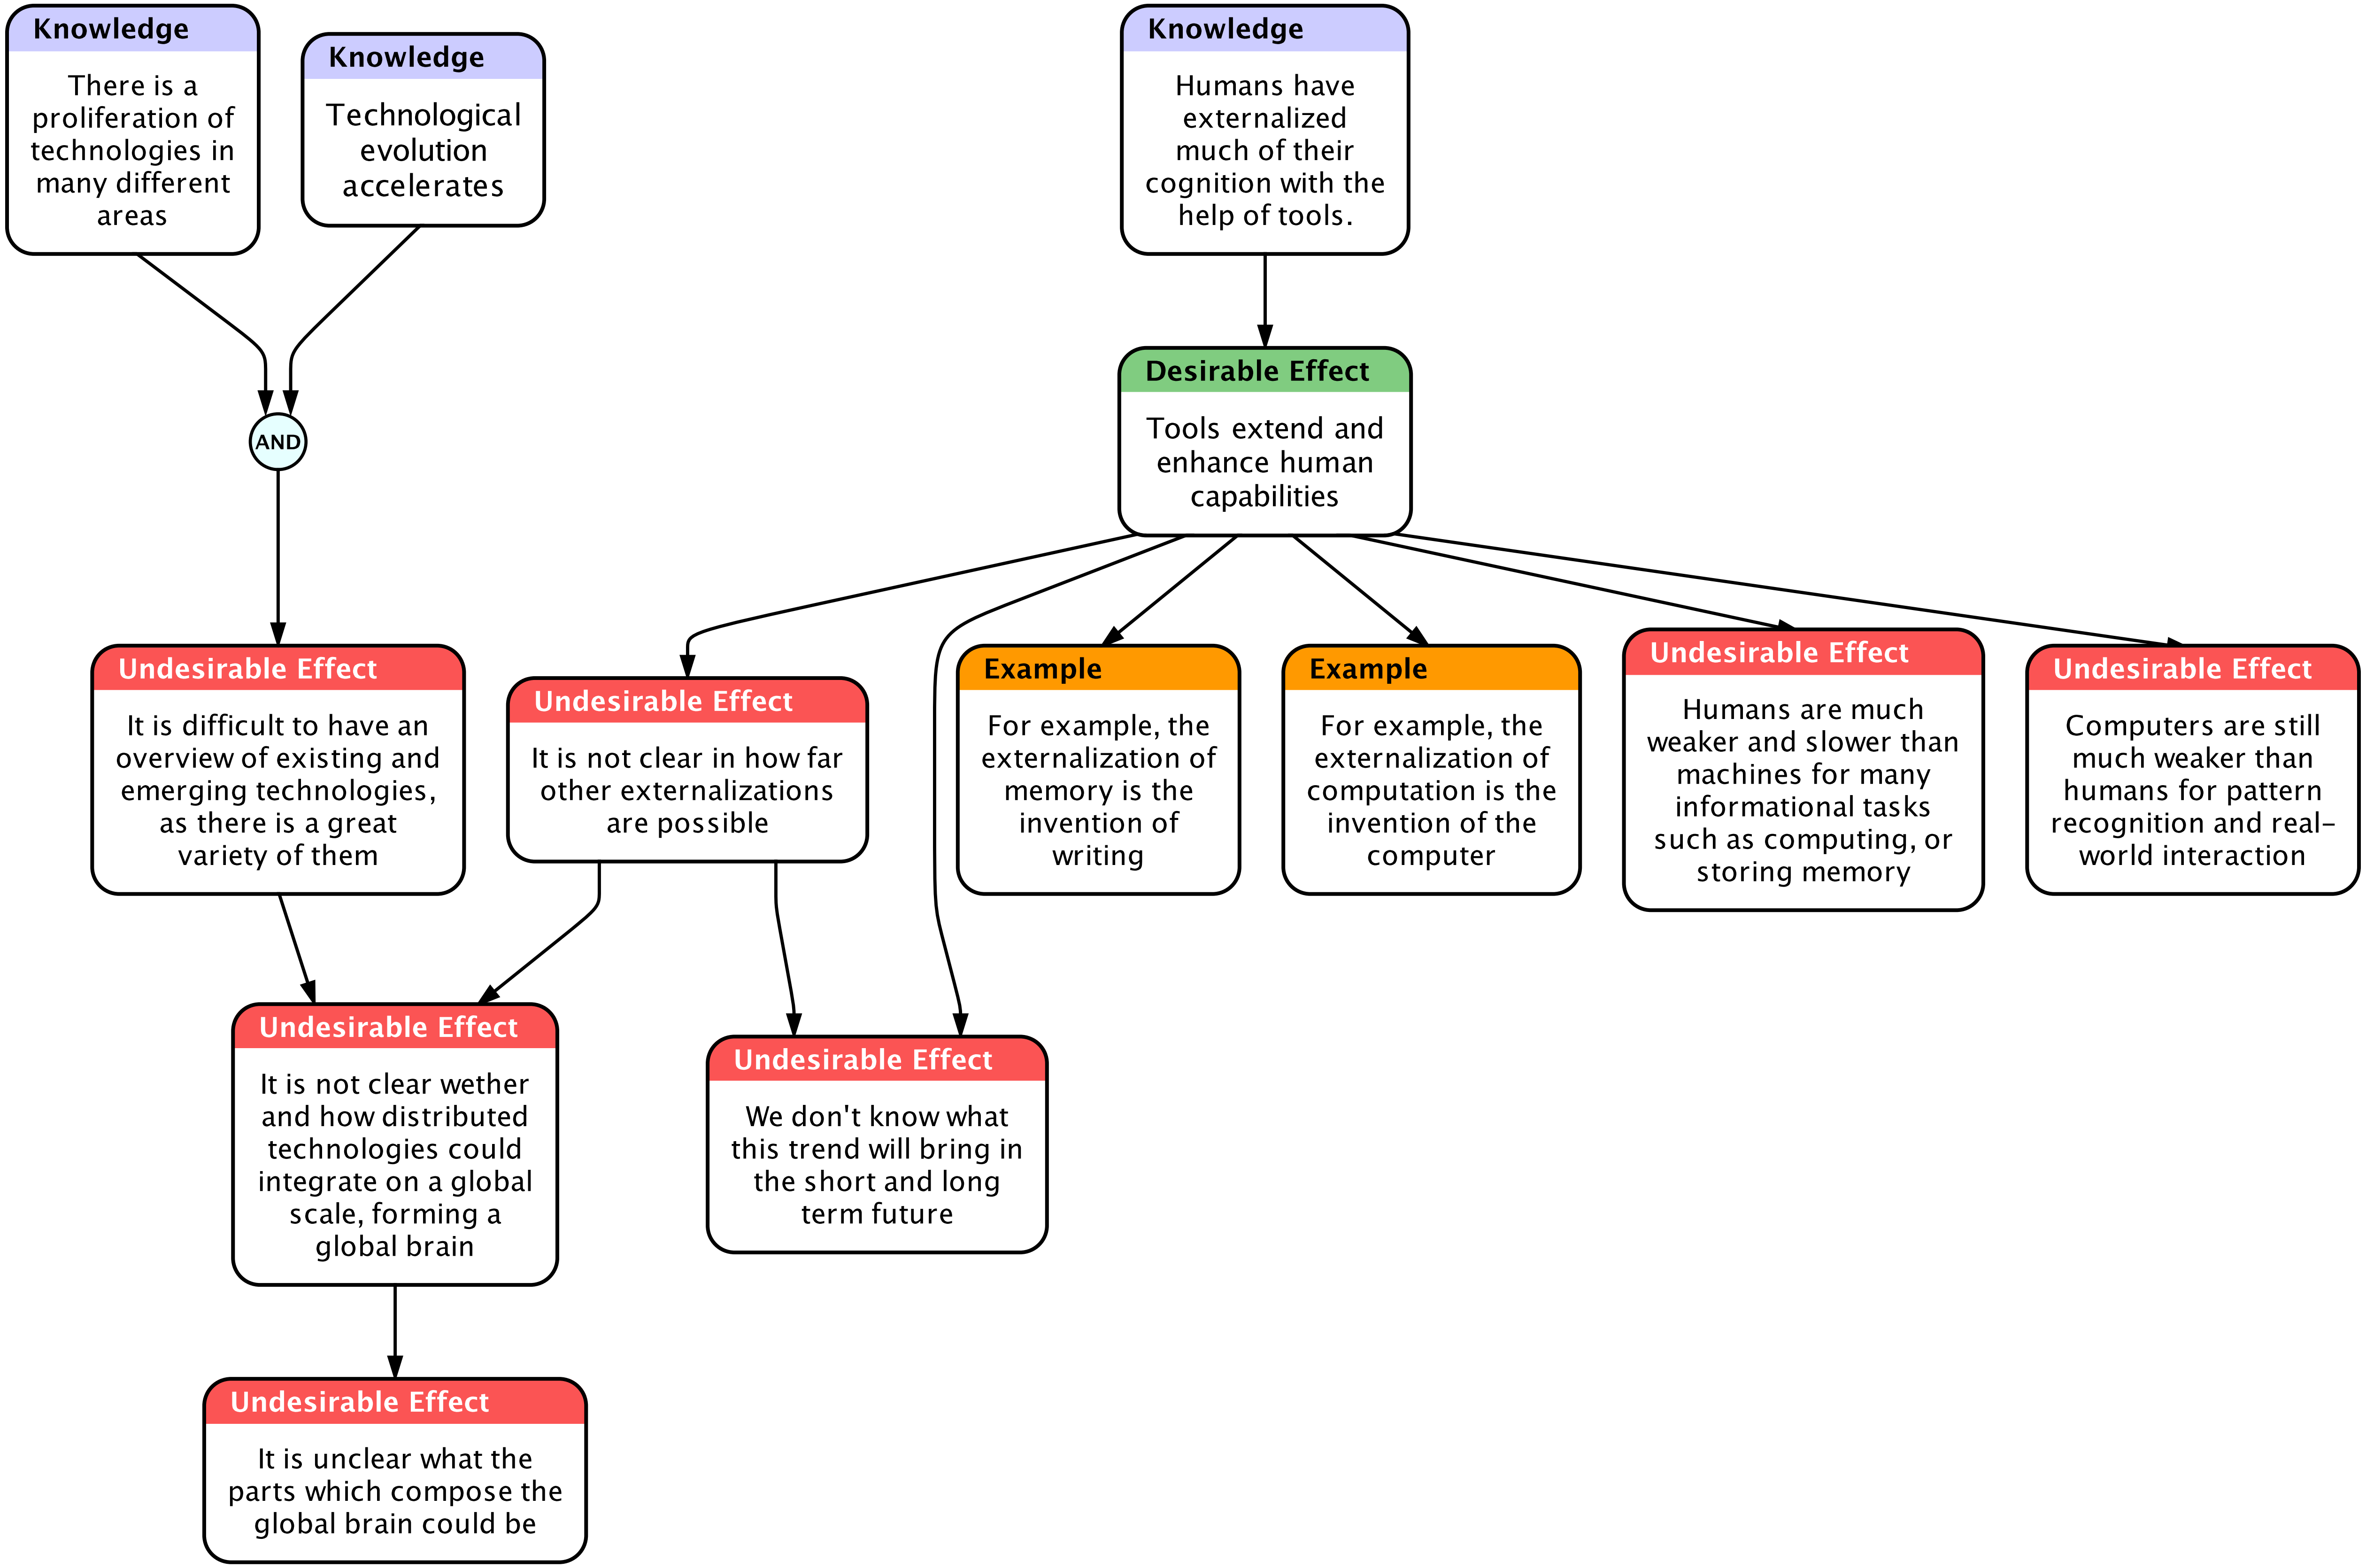
Knowledge
There is a proliferation of technologies in many different areas
Knowledge
Technological evolution accelerates
AND
Knowledge
Humans have externalized much of their cognition with the help of tools.
Desirable Effect
Tools extend and enhance human capabilities
Undesirable Effect
It is difficult to have an overview of existing and emerging technologies, as there is a great variety of them
Undesirable Effect
It is not clear in how far other externalizations are possible
Example
For example, the externalization of memory is the invention of writing
Example
For example, the externalization of computation is the invention of the computer
Undesirable Effect
Humans are much weaker and slower than machines for many informational tasks such as computing, or storing memory
Undesirable Effect
Computers are still much weaker than humans for pattern recognition and real-world interaction
Undesirable Effect
It is not clear wether and how distributed technologies could integrate on a global scale, forming a global brain
Undesirable Effect
We don't know what this trend will bring in the short and long term future
Undesirable Effect
It is unclear what the parts which compose the global brain could be

Section 1 - Local Brains or the externalized parts
Proposition
There are three phases in the externalization process
Proposition
The bodily self
Proposition
The extended self
Proposition
The networked self
Undesirable Effect
Is very limited
Intermediate Effect
Enhances the self
Intermediate Effect
Is enhanced both socially and technologically with social connections, as well as software and hardware agents
Knowledge
The scope of input, processing, output and feedback is restricted by bodily limitations
Knowledge
The scope of input, processing, output and feedback extends to the environment but remains local
Knowledge
The scope of input, processing, output and feedback is global
Desirable Effect
Externalizing cognition augments the scope and power of cognition
Proposition
We can externalize and distribute cognitive functions such as memory, computation, hearing, vision, brainstorming, reasoning, navigation, emotions and actions
Desirable Effect
Externalizing cognition enables sharing and distributing cognitive functions
Section 3 - The Global Brain or the emergent whole
Proposition
Depending on the kind of problem
Knowledge
We can use machines only
Knowledge
We can use humans only (crowdwork)
Knowledge
We can combine humans and machines, connected to the internet
Knowledge
This is the road of classical artificial intelligence
Proposition
The global brain vision proposes to integrate techno-sociological evolution into a globally distributed intelligence
Sub-issue
However, it is not certain that the human and machines combination is the most efficient one in the longer term
Hypothesis
The global brain parts could be humans, machines but also other biological cyborgs
Desirable Effect
Distributing cognition can be seen as an act of generosity leading to a new level of intelligence on planet Earth
Desirable Effect
The symbiosis of humans and machines is a framework to systematically explore, anticipate and integrate new technologies
Desirable Effect
Small entities which could sustain a wide range of temperatures would be fit for space exploration, high processing speed and intense energy usage
Section 4 - A Cosmic Perspective
Knowledge
Energy consumption growths exponentially
Knowledge
Memory, information processing and network bandwidth capacities grow superexponentially
AND
Indicator
An energy crisis is nearer than the singularity
Hypothesis
The starivore hypothesis suggests that other civilizations in the universe went over this crisis and now feed actively on stars
Desirable Effect
In the long term, if the starivore hypothesis is confirmed, it provides a glimpse of our far-future
AND
Desirable Effect
Both bottom-up and top-down analyses connect the dots between our present, our near future towards a global brain, and our far-future beyond the technological singularity.