

The GBI Vision: past, present and future context of global brain research

Francis Heylighen

Abstract: The following provides background and context for the creation of a “Global Brain Institute” (GBI), summarizing the scientific theories and research that will form the focus of this new institute. After defining the subject as the emergence of distributed intelligence supported by the Internet, it quickly summarizes the history of research in this domain, and then focuses on the present situation as investigated by the ECCO research group. The document suggests that the impact of this research remains limited because of lack of funding and infrastructure, and proposes to solve this problem by creating a Global Brain Institute at the Vrije Universiteit Brussel. The text concludes by sketching conceptual foundations for a future theory of the global brain: the global brain would emerge by self-organization in order to coordinate human and machine activities via the directed propagation of challenges from agent to agent across a global medium. It would thus play the role of a nervous system for the planetary organism.

Motivation

Tim Berners-Lee's breakthrough invention of the Web stems from a simple and easy way to link any kind of information, anywhere on Earth. Since then, the development of the web has been largely an erratic proliferation of mutually incompatible Web 2.0 technologies with no clear direction. This demands a new unified paradigm to facilitate their integration. The present document proposes to develop a theory of the global brain that would help us to understand and steer this on-going evolution towards ever-stronger interconnection between humans and machines. If successful, this would help us achieve a much higher level of distributed intelligence, thus allowing us to efficiently tackle global problems too complex for present approaches.

Definition of the Global Brain

The *global brain* (GB) is the idea that all people of this planet together with the information and communication technologies that connect them form the equivalent of a nervous system or “brain” for the planet Earth [1], [2]. While this idea may initially have been formulated merely as a metaphor, globalization together with the explosive development of the world-wide web are turning it into an increasingly realistic model of the present information society. Indeed, the growing network of social, communication, and economic links make individuals, organizations, machines and even ecosystems across the world ever more interdependent. Moreover, the storage, exchange and propagation of information across this networks provides it with a level of knowledge and capability for intelligence that far surpasses that of any individual or organization. An extrapolation of these accelerating changes trends points towards a “metasystem transition” or “singularity” within the next few decades [3], [4]. This is the emergence of a new level of distributed intelligence that is as yet difficult to imagine, but that will radically alter the human condition.

A short history of the GB vision

In the late 19th century, the founding fathers of sociology, Émile Durkheim and Herbert Spencer [5], observed that society is in many aspects similar to an organism, but did not yet find any clear counterpart for a nervous system in this *social organism*. In the 1920s, the French paleontologist Teilhard de Chardin [6] described the growth of the *noosphere*, the network of ideas and communications that envelops the planet, a concept he developed together with the Russian geologist Vladimir Vernadsky [7].

Teilhard's very abstract vision was complemented by the more practical approaches of the Belgian information scientist Paul Otlet [8], who envisaged a world-wide web-like interface to the whole of human knowledge, and of the British author H. G. Wells [8], who pleaded for the creation of a "World Brain", by which he meant a university-like global institution that would collect, organize and make available all that knowledge. For a practical implementation of these plans, we had to wait for the technology of hypermedia further developed by the Americans Vannevar Bush [9], Douglas Englebart [10] and Ted Nelson [11], and the emergence of the Internet in the 1970s. Internet and hypermedia were first integrated by the British computer scientist Tim Berners-Lee [12], who thus in 1991 created the *World-Wide Web*, an invention that would soon take over the world.

While the web, with its network of associative links, was clearly inspired by the organization of the brain, the link with the social organism was still lacking. This link was clarified by a number of authors inspired by Teilhard's vision: the British physicist Peter Russell [13], who coined the term "global brain" in 1982, the German complexity scientist Gottfried Mayer-Kress [14], who connected Russell's idea with the Internet, the French futurist Joël de Rosnay [15], [16],[17], who discussed the "planetary brain" of the "global macro-organism", and especially the Russian physicist and computer scientist Valentin Turchin [18]. As one of the founding fathers of Artificial Intelligence in the Soviet Union in the 1960s, Turchin developed an integrated theory of the evolution of cybernetic organization and intelligence, from primitive cells to the human brain, and beyond, to what he called the social "superbeing". His core innovation was the concept of *metasystem transition* [19],[20]: the evolutionary emergence of a higher level of complexity or intelligence. The implication of his theory was that humanity is at present undergoing a metasystem transition to a level of intelligence that we as yet cannot imagine.

After moving to the USA, Turchin came in contact with the American cybernetician Cliff Joslyn, who suggested to collaboratively develop Turchin's ideas via the new tools of hypermedia and the Internet [21]. To do this, in 1989 they founded the *Principia Cybernetica Project* [22], being joined one year later by the Belgian cybernetician Francis Heylighen. Heylighen was quick to realize the importance of the new world-wide web to realize Joslyn's vision. He therefore created the Principia Cybernetica website in 1993 [23], as one of the first complex, collaborative websites in the world—which is still quite influential. While working with his then PhD student, the Belgian psychologist Johan Bollen, Heylighen further realized that the world-wide web could be made much more intelligent, by implementing the mechanisms of Hebbian learning and spreading activation that characterize the brain. Combining these insights with Turchin's theory led him to propose a first concrete model of the future, intelligent web, i.e. the global brain, in 1996 [24].

After coming into contact with the American artificial intelligence researcher Ben Goertzel, who had developed similar notions [25], the two of them founded the international *Global Brain Group*. This brought together most of the researchers who had actively reflected about this issue, including Russell, Mayer-Kress, de Rosnay, Turchin, Joslyn and Bollen. The group organized a first international workshop on the global brain in 2001 [26] at the Free University of Brussels (VUB), in which some twenty researchers presented their views, and has since maintained an active email discussion forum on the topic.

Present situation

Over the past decade, both the Principia Cybernetica Project and the Global Brain Group have remained largely non-active—in part because of the illness and eventual death of the founding members Turchin and Mayer-Kress. Moreover, most of the other members became increasingly busy with more specialized activities which left them little time to reflect about the general topic of the next metasystem transition.

This fragmentation and hectic activity was also typical of the Internet as a whole over that period: millions of people flocked to the web in order to develop an enormous variety of websites, communities and software systems, which unfortunately were largely incompatible, lacking the simple, unified vision that had originally driven Berners-Lee. The dominant paradigm in that decade was the so-called “Web 2.0”, which emphasized the social, collaborative aspects of a web to which everyone could contribute—without realizing that this democratic process was already part of the original “Web 1.0” vision of pioneers such as Berners-Lee, Nelson, Englebart, and Joslyn, with the result that the wheel was constantly being reinvented.

Most recently, this chaotic mixture of successful and (mostly) failed experiments seems again to be moving towards some form of consolidation around a number of big players, such as Google, Apple, Wikipedia, Facebook and some large open-source communities. This prepares the minds for another look at the wider picture of where the Internet is heading in the medium and long term. This feeling was recently expressed by Tim O’Reilly, the Web 2.0 guru, in a keynote talk on the emerging global brain [27]—which, however, seemed to ignore most of the research that had already been done on that topic...

In a sense, the situation nowadays is similar to the one on the Internet in 1990, just before Berners-Lee introduced the web as a simple unified interface and linking structure for the many incompatible systems (ftp, gopher, telnet, newsgroups, WAIS, email...) that then existed on the Internet. The proliferation of mutually incoherent Web 2.0 systems demands a unified paradigm that would seamlessly integrate them all. However, the situation is much more complex than in 1990 because both the number of users and the number of uses of the Internet continues to explode. Just agreeing about common standards and protocols, as the Semantic Web initiative [12] has been trying to do, is not sufficient to deal with this burgeoning complexity. Instead we need a theory that takes into account the self-organizing, evolving and to some degree chaotic nature of such a complex social and cognitive system. The global brain concept promises to do just that.

Research in ECCO

In recent years, Heylighen has assembled a broad, international group of researchers around him at the VUB, in essence to extend the evolutionary-cybernetic theory pioneered by Turchin and the Principia Cybernetica Project. This “Evolution, Complexity and Cognition” (ECCO) group has made important advances towards a general theory of the self-organization of collective intelligence [28]. The new theory sees a metasystem transition as the coordination of the actions of initially autonomous agents. This coordination takes place spontaneously via the mechanisms of stigmergy, Hebbian learning, and the propagation of challenges. The same mechanisms explain the spectacular effectiveness of the human brain, and of insect societies, which are capable of surprisingly intelligent behavior while consisting of individually stupid members. Partial implementations of these mechanisms can be found in the most successful Web 2.0 ventures, such as Wikipedia, the open-source communities, and various social media [29]. However, this seems to have happened more by chance than by conscious design.

In addition, over the past decade the members of the ECCO group (here including former member Johan Bollen, now at Indiana University, with whom ECCO maintains close contacts) have proposed and implemented a variety of algorithms and demonstrations of collectively intelligent systems inspired by the Global Brain paradigm. These include:

- a flexible system of algorithms to automatically learn the preferences of users and recommend the best options for them personally, based on Hebbian learning and spreading activation [30], [30];
- SNURF, a proposal for the intelligent propagation of messages across a social network, so that people would only receive the messages that are most interesting for them personally [31];
- Smartocracy, a prototype of a self-propagating voting network for intelligent, democratic decision-making [32];
- Gnowsys, a collectively editable conceptual network for knowledge organization and education [33][34];
- Teacherplex, a website that stimulates students to learn new material via a system of multiple choice questions that constantly challenges them and provides feedback about their progress [35][36];
- a high-resolution map of scientific knowledge based on the paths that users follow collectively while searching for information [37];
- a mathematical formalism and set of algorithms for intelligent reasoning and general computation across distributed knowledge networks [38];
- Stigmergic University [28], a proposal to integrate these different technologies and more in order to create an open, global system of collaborative knowledge creation (academic research), learning, and assessment of progress. The intention is to mobilize and coordinate its users to develop and disseminate knowledge in the most efficient way, thus implementing a “World Brain” as envisaged by H.G. Wells [39];
- a collective intelligence system that successfully predicts the movements of the stock exchange by analyzing the mood of messages on Twitter [40].

Impact of Global Brain Research

While the ECCO group and the associated Global Brain community have built up an impressive array of expertise, ideas and insights around the GB concept, the impact of their research is as yet still limited. The top authors in the ECCO group (Heylighen, Bollen, Gershenson, and Rodriguez) together have gathered nearly 6,000 scientific citations (3,300 for Heylighen alone) for their work, as measured via Google Scholar. This number is steadily increasing. However, it is still relatively little compared to the best-known authors in this domain (e.g. Tim Berners-Lee has 32,000 citations, Vannevar Bush 8,000; Teilhard de Chardin 6,800), albeit better than most of the unduly forgotten pioneers (e.g. Turchin: 1,500; Otlet: 630; Englebart: 140). Moreover, apart from a few popularizing articles in the mass media (including *New Scientist*, *Le Monde*, and the *Washington Post*), their ideas have as yet hardly spread beyond the academic community. In particular, they do not seem to have made any observable impact on the large-scale commercial or public exploitation of the web.

The major reason for this limited impact is the lack of direct funding for global brain research. The members of ECCO have largely been working on these topics as volunteers, while being paid for other work or even sometimes living from their own savings. It is difficult to find money for this research because: (a) government-supported science foundations do not recognize the global brain as a legitimate scientific subject, because it cannot be fitted into one of the traditional disciplinary subdivisions, such as computer science, sociology, or economics; (b) commercial enterprises are typically interested only in research that promises concrete applications in the short term, not in understanding the wider, long-term picture.

More generally, most people, who ignore the existence of both theoretical foundations and empirical support for the idea, tend to dismiss the concept of the global brain as either a mere metaphor for the existing web, or a highly speculative, futuristic vision that would rather belong in a science fiction novel than in a scientific study. Even worse, many people when hearing the idea make

unfortunate associations with collectivist, totalitarian systems that suppress individual freedom. Yet, a deeper analysis shows that the opposite is true [41]: a global brain would promote freedom, democracy, diversity and emancipation, rather than force individuals to conform to some higher authority.

For the global brain paradigm to get the impact it deserves, we clearly need to reach a critical mass of active researchers, publications, demonstrations, applications, and dissemination towards the wider public. This requires a substantial injection of money into the community. The most cost-effective way to do this appears to be the creation of a Global Brain Institute (GBI), which would employ or support the most advanced researchers in the domain, giving them the opportunity and the infrastructure to work full-time on developing the theory, applications, and empirical evidence. Moreover, such an institute would provide a highly visible center or forum, where people interested in these ideas can get informed and connect, and from which ideas can disseminate.

A list of concrete objectives for organizing such a GBI is proposed in the accompanying GBI Strategic Plan. They include the production of a substantial number of publications, talks, conferences, trainings, meetings, seminars, web pages, discussion forums, software prototypes, and in-depth studies. Most important, however, would be an integrated, comprehensive theory of the Global Brain and its dynamics, which should be formulated at least partially in a mathematical and algorithmic form. This encompassing theory would also form the basis for a regularly updated survey of the state-of-the-art in global brain research, so that interested parties can get an as complete possible picture of the most important trends and developments, and what they mean for the short- and long-term future.

Why the VUB?

We propose to create the GBI at the Vrije Universiteit Brussel, i.e. the (Dutch-speaking) Free University of Brussels, in Belgium. The first reason is, of course, that this is the place where Francis Heylighen and the ECCO group, which have done most work on the global brain, are presently located. It is also the place where the Principia Cybernetica Web has been maintained since 1993, and where the Global Brain Group organized the first (and as yet still only) conference on the global brain, in 2001.

This is not a coincidence. The “Free” in the name of the university refers to the *principle of free inquiry*, which implies the rejection of all dogmas, authorities, or a priori beliefs, and to the freedom from domination by any political, religious, or commercial interests. In practice, this philosophy means that the VUB authorities have always been very open-minded and quick to adopt new social and scientific developments—even if those were initially controversial. For example, the VUB was one of the first universities in the world to create a center for in-vitro fertilization—which would later discover the ICSI (sperm injection) technique that has allowed millions of babies to be born that could not have been conceived otherwise. It was also one of the first to create a truly transdisciplinary research center to integrate the social and the physical sciences (the Center Leo Apostel, with which ECCO is affiliated). Most recently, it created an interdisciplinary PhD degree that allows combining research from any of the disciplines, thus removing all administrative hurdles towards making a PhD on the subject of the global brain.

This pioneering spirit was particularly clear in the domain of ICT. The VUB was the first university in the country to acquire a variety of innovations that nowadays seem standard: a fully computerized library database, a degree in computer science, an Artificial Intelligence Lab (led by the highly cited Luc Steels), a campus-wide computer network, a free connection to the Internet for all researchers and students, and a website (the Principia Cybernetica Web). Nowadays, the university houses a variety of large research groups working on both technological and social aspects of the on-going ICT revolution (albeit mostly with a horizon much narrower than the one of the global brain).

Why Brussels?

Another reason why the VUB is an appropriate place for the location of a global brain institute is the city of Brussels. By several measures (e.g. number of headquarters of international organizations, percentage of inhabitants born in a different country, variety of languages spoken), Brussels is the most international or “globalized” city in the world. It is not only the capital of the European Union and of the officially tri-lingual Belgium, but also the headquarters of NATO and a variety of multinational companies. The result is that Brussels has a large population of “expats”, i.e. highly educated and internationally oriented, foreign-born professionals, together with a variety of specialized schools, services and companies catering to that population. It is also situated at the crossroads of Western Europe, having excellent transport, economic and political links with the rest of the world. In spite of this centrality, living and working in Brussels is still relatively relaxed and inexpensive, especially compared with much larger cities like Paris, London, Moscow, Tokyo or New York.

The result is that ECCO has always found it very easy to attract students, researchers and visitors from all around the world. Many scientists and professionals pass by in Brussels for a meeting or conference, or on their way to another city, and use the occasion to come and visit us at the VUB. The same will surely happen with a GBI. This accessibility and openness is very helpful for developing worldwide, informal networks and for attracting a culturally diverse community of contributors.

In the last few years the VUB has understood that, next to its tradition of innovative research, this global prominence of Brussels is the university’s major asset. It therefore has centered its growth strategy on “internationalization”, i.e. attracting an increasing number of academics and students from abroad, and getting involved in a large number of international networks and cooperation agreements. The creation of a Global Brain Institute fits perfectly in this strategy, and is therefore likely to get the full support from the academic authorities.

Foundations for a Global Brain theory

The most crucial part in the research we plan to do at the GBI is to develop a comprehensive theory of the future intelligent information society. This model should encompass the likely effects of the on-going ICT revolution at all levels: technological, economic, social, political, psychological, cultural, educational, medical, ecological, and even chemical and physical. While this may seem overambitious or even meaningless for someone thinking in terms of traditionally separate scientific disciplines, the perspective of systems theory and cybernetics from which most global brain research emanates provides a solid foundation for such transdisciplinary modeling[17], [42].

Self-organization of the global organism

The basic analogy that helps us to understand this complexity is to view the planet Earth as a “living system” [43] or “macro-organism” [17][41]. Such an organism has a *metabolism*, which is the network of reactions that distribute, store and process matter and energy, and a *nervous system*, which is the network of links that distribute, store and process information. The information is used to regulate the metabolism, and to plan, select, and coordinate the actions that exploit the necessary resources and cope with unavoidable disturbances. From this cybernetic perspective, the intelligence of an organism is not so much its capacity for abstract, symbolic reasoning, but its ability to smoothly coordinate its activity so that it can tackle various challenges in the most efficient way.

What is still lacking in this analogy is that as yet the planet cannot be seen as an integrated, coordinated system: the interactions between the different subsystems are still characterized by far too much confusion, conflict, waste and friction. Several GB theorists (e.g. [6][18]) proposed that the further evolution of humanity will produce a much more integrated system. This seems to agree with present observations: globalization makes people across the world ever more interdependent and cooperative, while reducing misunderstanding and conflict [4]. (For example, the percentage of people killed in wars and aggression has been steadily decreasing over the past half century). However, most of these observations remain vague about the mechanisms that produce integration and about their eventual outcome: what kind of organization will characterize a fully integrated global system?

The last two decades have seen great advances in the understanding of *self-organization* [44], the process by which a complex system spontaneously develops coordination between its initially independent parts. These advances come in part from the theories of complex adaptive systems, networks, and collective or distributed intelligence. ECCO has succeeded in synthesizing many of the most important ideas, thus producing a general theory of the spontaneous coordination of actions[45].

Coordination as the propagation of challenges

The basic idea of our theory is that any action performed by any agent (e.g. a person, a robot, or a computer) produces a *challenge* for one or more other agents [28]. A challenge is a stimulus, invitation or opportunity for performing a further action. This further action advances the situation created by the previous action. A challenge could be a message sent, a question asked, an answer provided, an edit made, or most generally any perceivable change that is practically meaningful for another agent. If the challenge elicits one or more new actions, these actions in turn are likely to create new challenges. For example, a message received may be forwarded to another person, produce a reply, or incite the receiver to check a webpage. The forwarded or replied message may induce another forward or reply, etc. Thus, challenges can be seen to *propagate*: they travel and spread from agent to agent, while changing in the process. The modification they undergo is normally intended to respond to the challenge, i.e. to improve the situation from the perspective of the agent that received the challenge.

However, what is improvement for one may appear as deterioration for someone else. If this happens, the agents are in a relation of conflict or *friction*, the one obstructing or working against the other. However, conflicts normally do not last, as actions and counteractions after many variations tend to settle on a solution that is acceptable for both parties [4]. This may be a compromise or, in the better case, a win-win situation that is beneficial to all.

This means that the overall dynamics of challenge propagation is one of long-term, global improvement. The mechanism is similar to the variation and natural selection of the “fittest” solutions that drives evolution. This self-organizing dynamics is in general non-linear, as small, initial changes may be amplified by positive feedback, while the eventual outcome (the “attractor” of the dynamics) tends to be stabilized by negative feedback.

An example can be found in Wikipedia, the web encyclopedia that everyone can edit [29]. A nearly empty page on an interesting subject constitutes a challenge to the reader: how can I improve this page? The more people contribute to the page, the more interesting the text becomes, and the more further readers it will attract. More readers means more people who are challenged to think of further improvements, and therefore more contributions. This is a positive feedback that leads to accelerating advance. The additions made by one person, however, may be deemed inaccurate by someone else, and therefore corrected (counteraction). The original author may disagree with the correction, and revert to the previous version (counteraction to the counteraction). However, after a number of back and forth corrections, normally a text will emerge that is acceptable to all. As the discussion of the topic converges to one that is deemed accurate and comprehensive by everyone, further additions and corrections become less frequent, and the text becomes relatively stable.

This example illustrates how collective intelligence emerges from the *propagation of challenges*: the different contributions become aggregated into a collective answer much better than any individual could have produced [45]. This Wikipedia example is simple, because here it does not really matter who does what and when. In more complex situations, the actions may need to be coordinated more tightly, using forms of organization such as alignment, synchronization, workflow, regulation and division of labor [45]. These different forms of coordination too can self-organize according to the same basic mechanism of propagation of challenges.

Improving the propagation medium

The crucial requirement for successful propagation is a medium in which challenges can be reliably registered and distributed [4]. In the above example, the Wikipedia website provides this medium. Without such a flexible and accessible medium, the collaborative development of such a gigantic encyclopedia by millions of contributors would simply have been impossible. The lesson we can draw is that the emergence of collective intelligence critically depends on the quality of the medium for propagation or communication. Therefore, the advent of a global brain may be accelerated by improving the propagation quality of the available media. The Internet itself is the best illustration of this principle: it has spectacularly facilitated worldwide integration by its qualities of being digital (meaning that information can be propagated without loss), nearly instantaneous in its transmission, globally distributed, and robust against a large variety of malfunctions.

Another important factor is the *directionality* of propagation: if everything would be propagated to everybody, people would drown in an ocean of messages, most of which are irrelevant—or at least of very low priority with respect to their concerns. This problem is avoided in Wikipedia by the fact that readers actively choose the pages they want to read, and ignore the rest. The disadvantage, however, is that they may never find a page that is very relevant to their interests because they simply don't know it exists. An alternative propagation method can be found in social networks, where people indicate which other people they know and trust, so that they get automatically informed about whatever those “friends” or “connections” consider interesting. While faster and more selective than Wikipedia, this is still a rather inefficient and passive system for propagating challenges.

Our SNURF proposal [31] envisages using such a social network for a more active and intelligent propagation: messages would be passed on automatically with more than one step at a time, while being selectively routed depending on the interest profile of the people. For example, a report about a football game that you liked may be propagated to a “friend of a friend” interested in sports, but not to one interested only in science and culture. Our particle grammars [38] provide a similar method for selectively propagating queries (here called “particles”), not to people, but to knowledge containers distributed across a semantic network. This in principle allows the network to perform general purpose computing and reasoning, thus creating a form of artificial intelligence that is truly distributed across the globe. The combination of both methods makes it possible to direct propagation across a network of both human and non-human agents.

The disadvantage of both social and semantic networks is that someone needs to explicitly indicate whether two nodes should be connected or not. In the human brain, networks are much more flexible: two neurons or concepts are connected via synapses that can vary in their strength. This link strength continuously adapts to experience, becoming stronger whenever the two nodes act “cooperatively”, becoming weaker otherwise. Thus the network can learn autonomously which node to connect to which other node using which link strength. As such, it self-organizes so as to optimize the strength and directionality of its linking patterns.

Such an adaptive network provides the perfect substrate for the propagation of thoughts [24], [4], or what is called at the neural level, “activation”. While thoughts propagate by following associative links, they become better, absorbing relevant pieces of knowledge here and there, while ejecting pieces that are incoherent with the stored knowledge. In a sense, each pattern of activation is a “challenge” for the part of the brain where it passes to improve that pattern. This propagation of

challenges continues until the thought seems ready to be converted into an external action (e.g. writing down a plan, or performing a complex movement).

This neural organization was the inspiration behind our original model [46] of the learning and thinking web. However, it has never been implemented on a large scale because websites at the time were too rigid to automatically change linking patterns. Present-day websites, such as Amazon, Facebook and YouTube have incorporated some of that “brain-like” linking ability, by suggesting associated items (books, videos, pages, people you may know...) via the technique of collaborative filtering or recommendation systems. The underlying mathematical methods are related to our learning web algorithms, as analyzed by Heylighen in 1999 [47]. However, the methods used by these websites, while difficult to examine because of commercial secrecy, may well be less powerful than the ones we have proposed [30], [48]. For example, when Facebook suggests “people you may know”, in practice only a small percentage of these people tend to be actual friends of yours. Yet, with all the data that Facebook collects about who knows who and who interacts with whom, it seems possible to make a much more accurate prediction about which people you may like to establish a connection with. Of course, this matter can only be decided experimentally, by actually using these data together with “GB-like” algorithms to make recommendations, and then see which percentage is accepted compared to the percentage accepted from existing recommendation algorithms.

This example illustrates a general principle: the intelligence of a GB depends critically on the precision with which challenges (e.g. “do you know this person?”) are propagated to the right agents. This precision is easy to measure: you just count how often the challenge is met successfully (e.g. “yes, this is a friend of mine”) as a percentage of the total number of challenges issued. This success rate can then be used as a feedback signal for the propagation system. Thus, it can change its own parameters so as to maximize its success rate [30]. Of course, what is counted as success will depend on the type of challenge and the type of agent, but in general it is easy to determine in how far an action was successful, if only by measuring the satisfaction of the user.

As another example, imagine that you would like to find the most efficient way to go from one place to another, taking into account present circumstances such as traffic density, pollution, availability of public transport, etc. This problem could be tackled by sending out a number of software agents (“particles”) that propagate across nodes representing the different modes of transport (roads, traffic lights [49], train schedules, bus routes, ...) until they reach the destination node, and then send back their estimate of the total journey duration, cost, or some combined efficiency measure. This would allow you to determine in real time the optimal path, and have it continuously updated as the circumstances change. Thus, the “directed propagation of challenges” paradigm is applicable equally well for social interactions, for quantitative optimization of physical processes, and for a combination of both. Thus, it may form the basis for an efficient nervous system that coordinates all activity in the global organism.

Improving how challenges are tackled

A third factor that determines the effectiveness of the propagation of challenges is the dimension of challenge itself: a challenge is a situation or message that stimulates an agent to act. That means that effective challenges must *motivate* their recipients to perform useful actions. The best ICT environments, such as the iPhone or certain games, are intrinsically stimulating or challenging; they provide a very pleasurable user experience that motivates the user to continue using the system. The simplest theory to explain such intrinsic motivation is Csikszentmihalyi’s concept of *flow*[50]. Flow characterizes the kind of intrinsically enjoyable activity in which you are so fully absorbed that you would not consider stopping it. Requirements for achieving flow are: 1) clear goals; 2) immediate feedback on the actions you perform; and 3) challenges in balance with your skills, i.e. neither too easy nor too difficult. A good coordination medium should automatically fulfill these requirements, e.g. by propagating clear challenges to the people with the best matching skills, and propagating the results of their interventions back to them.

In addition, psychologists and social scientists have uncovered a variety of other techniques that dependably increase readiness and motivation to act [51]. By combining such motivating principles with the general mechanism of coordination, we can turn an ICT environment into a *mobilization system* [28], i.e. a system that stimulates people to collaborate in the most efficient and gratifying manner towards a valuable end. Applied at the level of the Global Brain, such systems would not just enhance cognition and coordination, but motivation, unity of purpose, and the general feeling of satisfaction for the world population.

A final factor supporting the effective propagation of challenges is simply the availability of the right agents to deal with the challenge. Availability is not just a question of finding who is best suited to tackle the challenge, but making sure that that agent is at the right spot at the right time. Thanks to mobile Internet connected devices, people can be easily located or directed to the right spot if they are needed, for example to help with an emergency.

But many simple tasks are better tackled by hardware agents: sensors, thermostats, vehicles, machines, simple robots..., which can be positioned at the most appropriate spots to perform their functions. For example, simple temperature and pressure sensors can be spread out all across an environment, and propagate their measurements to a central weather monitoring and forecasting agent, thus allowing it to warn people in real-time about tornadoes or other localized events. Security sensors can be attached to doors of buildings, checking who enters and who leaves, and warning a control system of potentially suspicious behavior. Vice versa, the control system can send “passes” to particular security sensors, so that they know they should open the door to specific individuals when they recognize them (e.g. via an iris scan). Nearby robotic vehicles can be immediately dispatched to the place where an emergency call originated, so that they can gather more information and perhaps provide some initial help.

The broad distribution of a variety of such artificial agents leads us to the vision of “ambient intelligence”: our environment itself would start reacting intelligently to our needs and desires, by combining the distributed knowledge and intelligence of the global brain with the physical embodiment that is needed to sense and react to the local situation. The technology to build such sensors and effectors is relatively straightforward. More complex is the coordination between billions of such devices that are spread out all across the world. Here again, the theory of self-organization via the propagation of challenges can help us to ensure effective collaboration between these artifacts, as proposed by Gershenson and Heylighen [52][53]. One example worked out in detail by Gershenson are self-organizing traffic lights, which regulate traffic much more efficiently than centrally controlled ones by propagating synchronized “waves” of vehicles [49].

Towards an integrated theory

The previous considerations look like a solid yet flexible foundation for developing an integrated theory of the global brain—here viewed as a nervous system for the planetary superorganism. Such a nervous system would not merely distribute data, but messages (“challenges”) that are practically meaningful, in the sense that they direct and coordinate the actions performed by the superorganism’s components, i.e. the billions of human and machine agents. This coordination would optimize the general functioning of the superorganism in an objectively measurable way. That means that all the problems that nowadays are still characterized by a lot of needless trial-and-error, confusion, friction, stress and waste of resources would be tackled much more efficiently. This opens up as yet difficult to conceive perspectives in domains as diverse as knowledge creation, education, transport, governance, economics, sustainability, energy consumption, social interactions, work satisfaction, and general well-being for the world population.

The conceived theory would be both *descriptive* and *prescriptive*: (1) it would model existing trends in terms of their underlying dynamics of challenge propagation, self-organization and increase of synergy, thus allowing us to extrapolate these trends so as to forecast future developments; (2) it would suggest ways to accelerate and steer this on-going evolution, by focusing on the factors that crucially determine the efficiency of these processes of self-organizing coordination. This should

make it possible to design systems that propagate challenges much more effectively, e.g. by improving the underlying medium so that it registers relevant information more precisely, formulating better algorithms for the direction or “routing” of challenges, and ensuring that challenges meet the “flow” criteria, so that people would truly enjoy contributing to the collaborative work.

However, developing such a theory in all its concrete and formal details will demand a lot of effort. It requires at least a dedicated group of highly intelligent, open-minded and motivated researchers, with a variety of backgrounds, to investigate as many as possible of the different applications and aspects of the global brain theory. Global society in general, and the Internet in particular, are highly complex, diverse and quickly changing environments. Trying to model them as a whole without a good theoretical framework seems a hopeless enterprise. But even with a good conceptual foundation, as we presently appear to have, there is still a need to confront theory with practice, and to fill in the necessary concrete details.

About every day a new type of website, information technology, or social movement seems to be born. Some of these will turn out to change the world, as illustrated by recent phenomena such as Wikipedia, Twitter, Facebook, and the iPhone. Up to now, most of these seem to fit neatly into our “challenge propagation” framework, so that we have a basis to evaluate their importance and future potential. However, it is possible that a new movement, theory or technology will emerge that requires us to fundamentally revise our framework, or at least to extend it with some additional mechanisms. This implies a constant scrutiny of new trends and developments in a variety of domains. Only an extended, multidisciplinary team following all the major innovations can achieve this.

Another requirement for building a full theory of the global brain will be testing and experimenting by means of simulations. The challenge propagation framework, with its roots in the sciences of complex dynamic systems, self-organization, neural networks and cybernetics, in principle allows building mathematical models. Valentin Turchin already started to develop a mathematical foundation for what he called the “cybernetic ontology of actions” [54], which was a major inspiration for our own ontology of action and challenge [28]. However, while the foundational concepts may be formalized relatively simply, their application to more realistic models will be too complex for an analytic mathematical treatment. This problem is traditionally tackled by developing simplified, but still relatively complex, computer simulations, where the model can then be tested across a variety of parameter values, initial conditions, or random perturbations [55]. Ideally, these computer models should also be tested with real-world data, e.g. from collaborative websites such as Facebook. Such simulation and testing demands a lot of time invested by highly qualified researchers.

Conclusion

This document has sketched the context for the creation of a Global Brain Institute at the Vrije Universiteit Brussel (VUB). It first reviewed the historical research that led to the concept of a global brain. It then focused on the deepening, widening and concretization of that research, first in the Principia Cybernetica Project, then the Global Brain Group, and finally the Evolution, Complexity and Cognition (ECCO) group at the VUB. It concluded that while the results are rich and promising, they have as yet had little impact outside a limited circle of academic researchers. The reason is that these research groups were basically run on a volunteer basis, so that their members were not able to dedicate their full time and attention to the problem. Therefore, they were not as yet able to provide all the explicit mechanisms, implementations, and evidence that would convince an outsider of the importance of this research.

The text proposed to solve this problem by creating a dependably funded, interdisciplinary institute that is fully committed to developing, testing and disseminating a theory of the global brain.

It argued that Brussels in general and the VUB in particular are a good location for such an institute, because of their extensive international connections and tradition of unconstrained, innovative and interdisciplinary research.

The text concluded by providing a first exposition of the conceptual framework developed in ECCO as a foundation for an eventual global brain theory. It sees the global brain as a nervous system for the planetary organism, which functions to direct and coordinate all the activity performed by the organism's components—conceived here as (human and artificial) agents. This coordination emerges via the mechanism of self-organization, based on the propagation of challenges from agent to agent. Any medium that facilitates such directed propagation, such as the Internet, the web, Wikipedia, or Facebook, increases the ability of the agents to collaboratively solve problems, and therefore their collective intelligence. The discussion concluded with the need to develop a (preferably mathematical) model of challenge propagation, which would help us to understand, anticipate and optimize this process of distributed coordination, and therefore the development of a global brain. But this can only be achieved by a sufficiently large group of dedicated researchers with a variety of disciplinary backgrounds.

References

- [1] F. Heylighen, "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., Uchitel Publishing, 2011, pp. 274 - 289.
- [2] F. Heylighen, "The global brain as a new utopia," in *Zukunftsfiguren*, Suhrkamp, Frankfurt, 2002.
- [3] F. Heylighen, "A brain in a vat cannot break out: why the singularity must be extended, embedded and embodied," 2012.
- [4] F. Heylighen, "Accelerating socio-technological evolution: from ephemeralization and stigmergy to the global brain," in *Globalization as evolutionary process: modeling global change*, vol. 10, 2008, p. 284.
- [5] H. Spencer, *The principles of sociology*, vol. 1. D. Appleton and company, 1898.
- [6] P. Teilhard De Chardin, *The phenomenon of man*. Collins London, 1959.
- [7] V. I. Vernadsky, *The Biosphere*. New York: Copernicus, 1998, 1926.
- [8] P. Otlet, *Monde: Essai d'Universalisme*. Brussels: Mundaneum, 1935.
- [9] V. Bush, "As We May Think," *The Atlantic Monthly*, Jul. 1945.
- [10] D. Engelbart, "A conceptual framework for the augmentation of man's intellect," *Computer-supported cooperative work: a book of readings*, pp. 35–66, 1988.
- [11] T. Nelson, *Literary machines*. Swarthmore: The Author, 1983.
- [12] T. Berners-Lee and M. Fischetti, *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Harper San Francisco, 1999.
- [13] P. Russell, "The Global Brain: speculations on the evolutionary leap to planetary consciousness," *Los Angeles: JP Tarcher*, 1983.
- [14] G. Mayer-Kress and C. Barczys, "The global brain as an emergent structure from the Worldwide Computing Network, and its implications for modeling," *The information society*, vol. 11, no. 1, pp. 1–27, 1995.
- [15] J. De Rosnay, *Le cerveau planétaire*. Olivier Orban, 1986.
- [16] J. De Rosnay, *The Symbiotic Man: A new understanding of the organization of life and a vision of the future*. McGraw-Hill Companies, 2000.
- [17] J. De Rosnay, "The Macroscopic: a New World Scientific System," 1979.
- [18] V. Turchin, "The Phenomenon of Science. A Cybernetic Approach to Human Evolution," *New York: Columbia University*, 1977.
- [19] V. Turchin, "A dialogue on metasystem transition," *World Futures*, vol. 45, pp. 5–57, 1995.
- [20] F. Heylighen, "(Meta) Systems as Constraints on Variation: A classification and natural history of metasystem transitions," *World Futures* vol. 45, pp. 59-85, 1995.

- [21] F. Heylighen and C. Joslyn, "Electronic networking for philosophical development in the Principia Cybernetica Project," *Informatica*, vol. 17, no. 3, p. 285, 1993.
- [22] C. Joslyn, F. Heylighen, and V. Turchin, "Synopsis of the Principia Cybernetica Project," 1993.
- [23] V. F. Turchin, C. Joslyn, and F. Heylighen, "Principia Cybernetica Web," 1993. [Online]. Available: <http://pespmc1.vub.ac.be/>.
- [24] F. Heylighen and J. Bollen, "The World-Wide Web as a Super-Brain: from metaphor to model," in *Cybernetics and Systems '96. R. Trappl (Ed.). Austrian Society For Cybernetics*, 1996.
- [25] B. Goertzel, *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers, 2002.
- [26] F. Heylighen, "The First Global Brain Workshop: From Intelligent Networks to the Global Brain, Evolutionary Social Organization through Knowledge Technology," 2001. [Online]. Available: <http://pespmc1.vub.ac.be/Conf/GB-0.html>.
- [27] T. O'Reilly, *Tim O'Reilly on the Emerging Global Brain*. 2011.
- [28] F. Heylighen, "Self-organization of complex, intelligent systems: an action ontology for transdisciplinary integration," *Integral Review*, 2011.
- [29] F. Heylighen, "Why is Open Access Development so Successful? Stigmergic organization and the economics of information," *Arxiv preprint cs/0612071*, 2006.
- [30] F. Heylighen and J. Bollen, "Hebbian Algorithms for a Digital Library Recommendation System," Los Alamitos, CA, USA, 2002, p. 439.
- [31] M. A. Rodriguez, J. H. Watkins, D. J. Steinbock, and C. Gershenson, "Snurf: Social Networks for News Media Distribution," Los Alamos National Laboratory, LA-UR-06-2564, 2006.
- [32] M. A. Rodriguez et al., "Smartocracy: Social Networks for Collective Decision Making," Los Alamitos, CA, USA, 2007, p. 90b.
- [33] G. Nagarjuna, "GNOWSYS: A system for semantic computing," Tata Institute of Fundamental Research, 2005.
- [34] M. Kharatmal and G. Nagarjuna, "SELF Platform: A Teacher-Centric Collaborative Authoring System," *Journal of Applied Collaborative Systems Vol*, vol. 1, no. 2, 2009.
- [35] M. Martin, "Teacherplex: Enlightening the World," 2008. [Online]. Available: <http://www.teacherplex.com>.
- [36] M. Martin, "Teacherplex.com: A New Style of Website Devoted to Education," ECCO working paper 2009-06, 2009.
- [37] J. Bollen et al., "Clickstream Data Yields High-Resolution Maps of Science," *PLoS ONE*, vol. 4, p. e4803, 2009.
- [38] M. A. Rodriguez, "General-Purpose Computing on a Semantic Network Substrate," in *Emergent Web Intelligence: Advanced Semantic Technologies*, Y. Badr, R. Chbeir, A. Abraham, and A.-E. Hassanien, Eds. London: Springer London, 2010, pp. 57-102.
- [39] H. Wells, *World Brain*. Ayer Co Pub, 1937.
- [40] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *Journal of Computational Science*, vol. 2, no. 1, pp. 1-8, 2011.
- [41] F. Heylighen, "The Global Superorganism: an evolutionary-cybernetic model of the emerging network society," *Social Evolution & History*, vol. 6, no. 1, pp. 58-119, 2007.
- [42] K. E. Boulding, "General Systems Theory-The Skeleton of Science," *Management Science*, vol. 2, no. 3, pp. 197-208, 1956.
- [43] J. G. Miller, "Living systems," 1978.
- [44] F. Heylighen, "Complexity and Self-organization," *Encyclopedia of Library and Information Sciences*, eds. MJ Bates & MN Maack Taylor & Francis, 2008.
- [45] F. Heylighen, "Self-organization in Communicating Groups: the emergence of coordination, shared references and collective intelligence," in *Language and Complexity*, Barcelona University Press, 2011.
- [46] J. Bollen and F. Heylighen, "Dynamic and adaptive structuring of the World Wide Web based on user navigation patterns," in *Proceedings of the Flexible Hypertext Workshop (ACM)*, 1997.

- [47] F. Heylighen, "Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map," *Computational & Mathematical Organization Theory*, vol. 5, no. 3, pp. 253–280, 1999.
- [48] A. Heylighen, F. Heylighen, J. Bollen, and M. Casaer, "A distributed model for tacit design knowledge exchange," *SID2005*, 2005.
- [49] S. B. Cools, C. Gershenson, and B. D'Hooghe, "Self-organizing traffic lights: A realistic simulation," *Arxiv preprint nlin.AO/0610040*, 2006.
- [50] J. Nakamura and M. Csikszentmihalyi, "The concept of flow," *Handbook of positive psychology*, pp. 89–105, 2002.
- [51] R. H. Thaler and C. R. Sunstein, *Nudge: Improving decisions about health, wealth, and happiness*. Yale Univ Pr, 2008.
- [52] C. Gershenson and F. Heylighen, "Protocol requirements for self-organizing artifacts: Towards an ambient intelligence," *Arxiv preprint nlin/0404004*, 2004.
- [53] C. Gershenson, "Design and Control of Self-organizing Systems," 2007.
- [54] V. F. Turchin, "The Cybernetic Ontology of Action," *KYBERNETES*, vol. 22, pp. 10-10, 1993.
- [55] C. Gershenson, P. P. Gonzalez, and J. Negrete, "Thinking Adaptive: Towards a Behaviours Virtual Laboratory," *Arxiv preprint cs.AI/0211028*, 2002.