Heylighen F. (2002): "The Global Brain as a New Utopia", in: R. Maresch & F. Rötzer (eds.) *Zukunftsfiguren* (Suhrkamp, Frankfurt)

The Global Brain as a New Utopia

Francis Heylighen
Center "Leo Apostel", Free University of Brussels
http://pcp.vub.ac.be/HEYL.html

The role of utopias

Utopias seem to be out of fashion these days. The old ideologies, such as communism, have shown their inadequacies, and the "laissez-faire" liberalism that has replaced them is coming under more and more criticism. Instead, the intellectual climate has turned to either gloom and doom, or an "anything goes" postmodernist relativism. This general pessimism and apathy is fed by an increasing amount of bad news about wars, terrorism, ecological catastrophes, global warming, rising xenophobia and inequalities.

Yet, objective statistics show that humanity is doing rather well on a global scale: life expectancy, wealth, education level, democracy, safety, and even average IQ have been consistently rising for the world as a whole, and this for at least the past century [Heylighen & Bernheim, 2000a; Simon, 1995]. Still, it cannot be denied that accelerating changes in all segments of science, technology and society and growing complexity produce a lot of stress and uncertainty. While, on average, developments seem to be for the better, overall social evolution has become much more unstable, unpredictable, and uncontrollable [Heylighen, submitted], raising the risks of major catastrophes. The anxiety, distrust and hopelessnes that this engenders is made worse by the lack of a shared *world view* [Apostel et al. 1994, Heylighen & Bernheim, 2000b]. Such an encompassing philosophy would provide clear goals and values, and a positive vision of the future, motivating people to work towards a common project, and giving a meaning to their life.

The role of a utopia is to create such a vision. Past utopias have been discredited mainly because their views were unrealistic: they failed to take into account the complexity of social systems and the tensions and contradictions inherent in human nature. As a result, their application to real societies have all too often produced spectacular failures. Yet, recent developments such as evolutionary psychology and the theory of complex adaptive systems [Heylighen, 2002], are providing fundamental insights in these complexities, opening up the perspective of a utopia that takes them duely into account. While a realistic model of a future society will necessarily be complex, abstract and with many shades of grey, a utopia must in the first place provide a simple and appealing picture, that people can easily visualize. These two

requirements—of realistic complexity and of intuitive appeal—are inherently at odds, and therefore an acceptable utopia will have to be a compromise. In my own reflections on this issue [Heylighen, 2002], I have settled on the metaphor of the *"global brain"* as the for me most satisfactory vision of a positive, yet realistic future for humanity.

The global brain vision

The "global brain" is the name given to the emerging intelligent network that is formed by all people on this planet together with the computers, knowledge bases and communication links that connect them together. Like a human brain, this network is an immensely complex, self-organizing system, that processes information, makes decisions, solves problems, learns new connections and discovers new ideas. It plays the role of a nervous system for the whole of humanity. No single person, organization or computer is in control of this system: its "thought" processes are *distributed* over all its components. This concept of a cognitive system at the planetary level has been proposed by many different authors under different names: planetary brain, world brain, global mind, noosphere, social brain, Metaman [Stock, 1993], super-organism [Heylighen, 2002], super-being [Turchin, 1977], and collective consciousness are some of the roughly equivalent synonyms.

The history of this idea goes back quite a while. As the variety of names indicates, many people have independently developed the idea of society as an organism-like system with its own nervous system [Stock, 1993], each adding their own insights to our present understanding. The evolutionary theologist Teilhard de Chardin [1955] was probably the first to focus on the mental organization of this social organism, which he called the "noosphere". Around the same time, the science fiction writer H. G. Wells proposed the concept of a "world brain" as a unified system of knowledge, accessible to all. The term "global brain" seems to have been first used by Russell [1995]. The first people to have made the connection between this concept and the emerging Internet may well be Mayer-Kress [1995] and de Rosnay [2000]. Heylighen and Bollen [1996], and Goertzel [2001] appear to be the first researchers to have proposed concrete technologies that might turn the Internet into an intelligent, brain-like network.

The global brain can be conceived most fundamentally as a *higher level of evolution*, the way humans form a higher level of organization that evolved out of the animals. Although the analogy between an organism and a society can be applied even to primitive societies, it becomes clearly more applicable as technology develops. As transport and communication become more efficient, different parts of global society become more interdependent. At the same time, the variety of ideas, specializations, and subcultures increases. This simultaneous integration and differentiation creates an increasingly coherent system, functioning at a much higher level of complexity [Heylighen, 2002]. The emergence of such a higher order system has been called a "metasystem transition" [Turchin, 1977]. Examples of metasystem transitions include the origin of life and the development of multicellular organisms out of single celled ones. The appearance of a global brain, functioning at a much higher level of intelligence than its human components, seems a prime example of such a metasystem transition.

The global brain vision draws part of its inspiration from a number of related approaches. Collective intelligence [Lévy, 1997] is the idea that a group can be more intelligent than its members. The best known examples are social insects, such as ants, termites or bees, which are individually dumb, but capable of surprisingly intelligent behavior when functioning as a group. The intelligence of the global brain will be collective, as it arises from the interactions between millions of individuals. Symbiotic intelligence is the idea that intelligence can also emerge from the interactions between essentially different components, such as people and computers. As de Rosnay [2000] proposes, people will live in symbiosis with this surrounding network of technological systems, and out of this symbiosis, a higher level intelligence will emerge.

Gaia (the Greek goddess of the Earth) is the name given to the hypothesis that the planet Earth itself is a living organism. This organism would be able to regulate its own essential variables, such as temperature and composition of the atmosphere. Compared to the global brain as we have defined it, this "Gaian" organism seems very primitive, with a level of intelligence comparable perhaps to the one of a bacterium. At present, Gaia and the global brain are still largely independent, and the impact of society on the global ecosystem appears unsustainable. However, several authors have argued that Gaia and the superorganism will evolve to a state of symbiosis, that may eventually lead to a merging of the two. Thus, the global brain would not only form a brain for humanity, but for the whole of Planet Earth.

Although most researchers have addressed the global brain idea from a scientific or technological point of view, authors like Teilhard de Chardin [1955] and Russell [1995] have explored some of its spiritual aspects. Similar to many mystical traditions, the global brain idea holds the promise of a much enhanced level of consciousness and a state of deep synergy or union that encompasses humanity as a whole. Theists might view this state of holistic consciousness as a union with God. Humanists might see it as the creation, by humanity itself, of an entity with God-like powers. Followers of the Gaia hypothesis have suggested that the "living Earth" of which we are all part deserves awe and worship; it therefore could form the basis of a secular, ecologically inspired religion. The Global Brain vision may offer a similar sense of belonging to a larger whole and of an encompassing purpose.

Technologies for a Global Brain

In principle, it is possible to imagine a global brain even in the absence of information technology. As ideas ("memes") are communicated from person to person, they evolve, assimilating the contributions and points of views of myriad individuals. Thus, society already has a kind of a collective mind, constantly developing new thoughts that cannot be traced back to any individual contribution. However, in the absence of modern technology, this "collective thought" required decades to develop any new insights. Electronic media have made this process much more efficient, allowing ideas to spread and evolve in hours rather than decades. This turns the global brain from an interesting analogy into a phenomenon that can now be concretely experienced. Let me sketch how existing technologies can be further enhanced to produce a truly intelligent system.

The web is the hypermedia interface to the information residing on the Internet. It makes it possible to seamlessly integrate documents that are distributed over the entire planet, and created by people who may not even be aware of each others' existence. What holds these documents together is not their geographic location, but their *associations*: links connecting mutually relevant pages. This hypermedia architecture is analogous to the one of our brain, where concepts are connected by associations, and the corresponding assemblies of neurons by synapses. The web thus functions like a huge associative memory for society.

However, the brain is more than a static memory: it can *learn* and *think.* Learning takes place by the strengthening of associations that are used often, and the weakening of rarely used associations. Through learning, the brain constantly enhances its organization and increases its store of knowledge. Thinking happens by the activation of concepts and the "spreading" of this activation to related concepts, in proportion to the strength of association. Thinking allows the brain to solve problems, to make decisions, and to be creative, that is, discover combinations of concepts never encountered before. By making simple changes to its static architecture, we can implement similar processes on the web.

In the brain, learning follows the rule of Hebb: if two neurons are activated in close succession, the strength of their connection is increased. A similar procedure has been applied to the web by Heylighen and Bollen [1996]: if two web pages are consulted by the same user within a short interval, either the existing hyperlink between the pages gets a higher weight, or a new link is created. On any given page, only the links with a minimum weight are shown. Thus, links that are not sufficiently reinforced may eventually disappear. The result is that such a learning web constantly adapts to the way it is used, reorganizing its pattern of links to best reflect the preferences of its users. In practice, this creates direct links between the pages that are most strongly related, bypassing less interesting detours, and clustering pages together according to their mutual relevance. As such, the web becomes much more efficient to use, by assimilating the collective knowledge and desires of its users.

The simplest way to implement web "thinking" is to create a specialized software *agent.* This is a program that works as a "delegate" of its user, autonomously collecting information that is likely to be interesting to its user. The agent can learn the user's preferences simply by observing which pages the user actively uses, or it can receive specific instructions (e.g. keywords) from the user. Given that preference profile, the agent can locate pages that satisfy the profile, and then use "spreading activation" to find further, related documents. It does this by "activating" pages in proportion to their degree of interestingness, and then propagating that activation according to the hyperlinks and their weights as learned from other users. Thus, it can discover new documents, that may not contain any of the initially given keywords, but that are still highly relevant to the query. This is especially useful when the user cannot clearly formulate the query, but only has an intuitive feel for it.

With such technologies, the web would become a giant, collective brain, which you could consult at any moment to get an answer to your questions, however unusual or

vaguely formulated they may be. Its thought processes would always be ready to enhance and extend your own thinking. To fully harness the power of this global brain, it should be constantly available. The rapid spread of mobile communication already offers universal access to the web, wherever you are. Further miniaturization will lead to wearable computers, incorporated in your clothing, with images projected on your glasses. Automatic recognition of speech, gestures and even emotions will make communication with the web much easier and more intuitive. In the longer term, we can foresee direct connections between computer and brain, through neural interfaces. This would allow you to communicate with the global brain simply by thinking, having your thoughts immediately sensed, understood, and enhanced. Your thoughts could also be directly turned into actions, as when you use the global brain to order a pizza, get a taxi, or switch on the heating, so that it would be nice and warm by the time you come home.

All these technologies already exist, either as prototypes or as applications working in a more limited environment. What remains to be done is further streamlining, and above all integration into a single, coordinated whole.

Some illustrative applications

To make the myriad potential benefits of such a system more concrete, let us try to imagine how a global brain-like network could help tackle the problems of congestion, pollution and delay that hamper present-day transport.

Since Spring 2001, the Brussels company for public transport (www.mivb.be) has been offering the beginnings of an intelligent support system on the web. It allows you to enter your precise location and destination. The system then calculates the quickest combined connection, taking into account the time schedules of all forms of public transport available (tram, bus, subway) and the time you would need to walk to and from the different stopping places. Moreover, it provides you with detailed guidelines, such as: “take the second side street on the left, ..., take the subway at the station Delta, get off at, ..., until you reach the museum at 12.19”.

By adding a few recently developed technologies, we could take this several steps further. First, the system becomes truly useful only when it can be consulted in real time, i.e. while you are on the journey. This is already possible using a palmtop computer with a wireless Internet connection. It might further include a Global Positioning System, so that the system would know where you are, without need for you to enter an address. In this case, the guidelines could be constantly updated, so that if by error you walk down the wrong street, the system would recompute the optimal route and guide you interactively to your destination.

Second, instead of you having to buy tickets, payment could be done automatically, by transferring the correct amount of digital cash from your account to the one of the transport company the moment you finish each leg of the journey. This would imply the additional option of choosing the least expensive route rather than the quickest. This becomes especially interesting if the system would include different companies or forms of transport (e.g. taxis vs. buses) with different prices and different

transport offers. Depending on the route and the time of day, overall cost, comfort and duration would vary, but your personal agent would know your preferences, and negotiate with the transport system to find the option that best suits your situation and budget.

Such a system optimizes travelling for the individual, but this can be extended to the collective level. Imagine that all individuals would communicate with the transport system the way we sketched it. In that case the system would know precisely how many people travel between any two destinations at any time of the day. This would allow it to determine the most efficient way to organize the transport network. A truly interactive system would moreover be able to adapt routes in real time, depending on the demand. Roaming vans or buses would be directed by the system to the place where travellers are waiting, using the most efficient route that combines their various locations and destinations. Price and waiting time can be negotiated between the traveller's agent and the system. In that way, the "invisible hand" of the market (albeit electronically supported) can adjust supply and demand of transportation in the most efficient way.

At this level of flexibility, the distinction between public and private transport becomes moot: such a system can also accommodate individual drivers, who may be willing to share their car with passengers for part of their trip, in return for a fee negotiated between the respective agents. This is a more flexible version of car-pooling or hitchhiking. The proposed supply-and-demand driven merger of public and private transport does not imply that market forces should reign supreme, though. A truly *public* transport can reduce pollution, noise, energy consumption and congestion in a way that benefits society in the long term, without therefore having direct benefit to the traveller here and now.

Such collective benefits determine additional variables that should be taken into the equation when calculating the recommended route. They can be expressed as constraints that prohibit certain routes, e.g. a noisy bus passing near to a hospital. Others can be expressed as additional weights that bias the overall decision to the one that is optimal for society rather than for the individual traveller or transport firm. For example, the government can subsidize certain instances of transport and tax others, depending on their relative advantages and disadvantages. These corrections to the market price would be automatically taken into account when calculating the cost of the different options for the traveller. For example, the system may offer the train as cheapest option because it is less polluting, even though the immediate cost of a bus journey might be lower.

Such corrections should be able to adapt in real time, e.g. increasing the cost of more polluting options in proportion to the amount of air pollution that is registered by sensors. In that way, the present level of pollution can be regulated cybernetically. Moreover, the system would be able to react *immediately* to sudden perturbations, e.g. by redirecting traffic around the obstruction created by an accident. Thus, it would be able to prevent the self-reinforcing processes that lead to the build-up of a traffic jam,

and that in a sufficiently dense traffic can be triggered by a mere local slowdown involving one or two cars.

This transport example can be easily generalized to other problem domains. The generalization from transport of people to transport of goods is straightforward: it suffices to replace the location, destination and preferences of the traveller by those of the supplier and client of the good.

The generalization from transport to production requires a somewhat higher level of abstraction: every production process can be represented as a trajectory in *state space*, moving from the "raw" state of materials or components to the "finished" state via a number of intermediate states. This trajectory can be optimized taking into account the moment-to-moment supply of raw material and demand for finished goods. For example, when the electronic payment systems of booksellers register a higher demand for a particular novel, printing presses can be adjusted to immediately start producing more copies, while increasing the standing order for paper from the suppliers. Effects at the collective level (e.g. some processes are more polluting or consume more scarce resources than others) can be taken into account by imposing additional constraints or preferences.

Services, such as medical treatment, too can be conceptualized as a trajectory from an initial state (e.g. the patient being ill) to a desired state (e.g. the patient being cured). Again, an intelligent system can try to find the "shortest" route from the one state to the other, taking into account various constraints (e.g. avoiding risk to the patient's life) and optimization criteria (e.g. minimizing cost, pain, and side-effects), at both the individual and collective level, while being guided by the collective experience of other patients with similar ailments as stored in the shared memory.

General benefits to be expected

Now that we have a better grasp of how a global brain-like system would function in practice, let us try to summarize its great advantages for society.

A first problem that it would tackle is information overload [Heylighen, submitted]. As more and more documents, services, and people move to the Internet, retrieving, sending and receiving information becomes in practice effortless and free. This means that ever larger amounts of potentially interesting messages, documents and announcements will clamor for our attention. However, attention, unlike information, will remain scarce: our brain is simply unable to attend to more than a few dozen messages a day. Therefore, we increasingly need support from a system that is capable to sort through billions of information items and select those that are most relevant to our particular situation and interests.

Let us then summarize its effects on the economy. The market is the collective system of transactions that helps supply to match demand, and thus to fulfill the need of the collective customer for products and services. A traditional market is rather inefficient, requiring a huge infrastructure of middlemen, specialized organizations such as stock exchanges and auctions, and communication channels. The Internet already

allows such transactions to take place much more quickly and transparently, with less cost and effort. This strongly reduces "friction", making the economy more efficient so that demand can be satisfied more rapidly, more accurately, and at a lower cost [Heylighen, 2002]. The global brain will not only facilitate direct communication between buyers and sellers, but help buyers to find the best value (e.g. through shopping agents to compare prices), and help sellers to get the best price (e.g. through automatic auctioning systems). The net effect is that growth increases, while inflation and economic instability decrease. Moreover, there will be less waste because of unsold items or goods shipped far away when there is demand around the corner. The direct incorporation of collective effects ("externalities") in the decision-making process, as illustrated in the transport example, will moreover allow a more efficient governance over the economy, thus protecting employees and consumers while reducing inequalities and pollution, without the added complexity, bureaucracy and rigidity that tend to accompany such interventions in a traditional political system.

The global brain will moreover help eliminate conflicts. It in principle provides a universal channel through which people from all countries, languages and cultures of this world can communicate. This will make it easier to reduce mutual ignorance and misunderstandings, or discuss and resolve differences of opinion. The greater ease with which good ideas can spread over the whole planet will make it easier to reach global consensus about issues that concern everybody. The free flow of information will make it more difficult for authoritarian regimes to plan suppression or war. The growing interdependence will stimulate collaboration, while making war more difficult. The more efficient economy will indirectly reduce the threat of conflict, since there will be less competition for scarce resources.

Of course, technology alone will not solve all the problems that threaten our planet: in the end, people will have to agree about concrete policies to tackle e.g. global warming or poverty. Yet, the global brain can support not only the process of reaching consensus on a plan of action, but also its practical implementation. For example, combating infectious diseases or pollution will require extensive monitoring of the number of infections or concentration of polluting gases in different regions. Information collected by local observers or by electronic sensors can directly enter the global brain, be processed to reveal underlying trends, and be forwarded to the people or institutions responsible for taking direct action.

The ultimate test of a utopia is whether it will make people happy. Statistics about life satisfaction in different countries have shown that people are more happy when their society provides them with greater health, wealth, safety, knowledge, freedom and equality [Heylighen & Bernheim, 2000a]. The global brain will greatly contribute to each of these fundamental values. The global brain itself will provide universal access to all of humanity's knowledge, and thus indirectly increase people's freedom to choose their own path, while providing them with more equal opportunities. Its effect on the economy will directly create more wealth, and indirectly resources to invest in medical care, education, safety measures, etc. Its support for the creation of new knowledge will

boost science and technology, and thus help them to solve the most pressing medical, social and ecological problems in a much shorter time span.

Utopia or distopia: should we be afraid of the global brain?

Before concluding we must address a number of recurrent worries and misunderstandings about the global brain ideal [Heylighen, 2002]. The use of terms like "collective intelligence" or "superorganism" tends to invite the comparison of the global brain with a "hive mind", that is, a collective in which the members all think and behave the same, lacking any autonomy or personal identity. This frightening prospect is most vividly illustrated by the "Borg", the race of cyborgs imagined by the creators of the science fiction series "Star Trek". The global brain, on the contrary, derives its intelligence precisely from the diversity and autonomy of the people that take part in it. If everybody would make the same choices, then the global brain would not be smarter than a single individual. It is because different people have different points of view and different experiences that together they can tackle more complex problems.

Another recurrent worry is privacy, and the fear the the global brain might become a high-tech version of "Big Brother". Since the global brain becomes more effective by monitoring user behavior, it may seem that it will get to know everything an individual has done on the net, including actions that this user would rather keep private, such as visiting a pornographic site. However, web learning algorithms do not need to know individual activities. They only need to know the collective frequencies of certain actions, e.g. that most users who visited the "Playboy" site also visited the "Penthouse" site, independently of who these users are. In spite of this anonymity of web learning, the algorithms can still provide personalized recommendations based on the user's preference profile, but this profile is kept strictly on the user's own computer, far from the prying eyes of others.

Another common fear is that belonging to an encompassing, collective system, such as the global brain, will limit individual freedom. It is true that for a global brain to be effective, the people participating in it will need to agree about a minimum set of common standards or rules to facilititate communication and cooperation. However, well-chosen rules will increase rather than decrease freedom. This can be explained by an analogy with the traffic code. Without such traffic rules as the obligation to drive on the right hand side of the road (or left, in Great Britain), traffic would be much more dangerous and more easily obstructed, effectively decreasing your freedom to jump into your car and drive wherever the roads may lead you. Since the same rules will democratically apply to everyone, the net result will be that dominant organizations, governments, or corporations will have less power to censor or impose their rules on the people who use the net. This loss of power will understandably be resented by these organizations, but should be welcomed by individuals as it will increase their freedom and autonomy.

Since Frankenstein and the Golem, many publications have warned us that the artificial creatures we are creating may escape our control and eventually take over. The global brain, however, is not intended to replace humanity, but to complement or

augment it. The global brain is controlled by all the people that are part of it. It is not an autonomous system that could suddenly decide no longer to obey commands. The global brain's intelligence, indeed its "mind" or "personality", emerges from the actions of all people collectively. If the people would decide no longer to use the network, then the global brain would simply stop to exist.

A final worry is that advanced information technology may increase the gap between haves and have-nots, and more particularly between those that have access to information and those that haven't. Although global brain technologies will be adopted most quickly by the wealthiest and best educated part of the population, this won't stop the underdogs from joining a little later. Internet technologies are quite inexpensive to install, compared to e.g. roads, electricity or running water. Moreover, as the global brain becomes more intelligent, it will become ever easier to use, requiring an ever lower education level for entry. Speech technologies will soon make the web available even for illiterates, and may teach them to read and write in the process. Thus, the global brain is a cheap and efficient way to increase the education level, access to information, and economic competitiveness in all regions of the world, helping Third World countries to bridge the gap with the wealthiest countries.

References

Aerts, D., Apostel L., De Moor B., Hellemans S., Maex E., Van Belle H., Van Der Veken J. (1994): Worldviews: From Fragmentation to Integration, (VUB Press, Brussels).

de Rosnay, J. (2000) *The Symbiotic Man*. McGraw-Hill.

Goertzel, B. (2001) *Creating Internet Intelligence: Wild Computing, Distributed Digital Consciousness, and the Emerging Global Brain* (Plenum)

Heylighen F. & Bernheim J.(2000a): "Global Progress I: empirical evidence for increasing quality of life", *Journal of Happiness Studies* 1 (3), p. 323-349.

Heylighen F. & Bernheim J. (2000b): "Global Progress II: evolutionary mechanisms and their side-effects", *Journal of Happiness Studies 1(3)*, p. 351-374.

Heylighen F. & Bollen J. (1996) “The World-Wide Web as a Super-Brain: from metaphor to model”, in: *Cybernetics and Systems '96* R. Trappl (ed.), (Austrian Society for Cybernetics), p. 917-922.

Heylighen F. (2002): "The Global Superorganism: an evolutionary-cybernetic model of the emerging network society", *Journal of Social and Evolutionary Systems*

Heylighen F. (submitted): "Complexity and Information Overload in Society: why increasing efficiency leads to decreasing control”

Lévy P. (1997), *Collective Intelligence: Mankind's Emerging World in Cyberspace* (Plenum, New York)

Mayer-Kress, G. & Barczys C. (1995) “The Global Brain as an Emergent Structure from the Worldwide Computing Network, and its Implications for Modelling”, *The Information Society 11:1*, 1-28.

Russell, P. (1995) *The Global Brain Awakens:*, Miles River Press.

Stock, G. (1993) *Metaman: the merging of humans and machines into a global superorganism.* Simon & Schuster, New York.

Teilhard de Chardin, P. (1955) *Le Phénomène Humain.* Seuil, Paris.

Simon J. L. (ed.) (1995) *The State of Humanity*. Blackwell, Oxford.

Turchin, V. (1977) *The Phenomenon of Science. A Cybernetic Approach to Human Evolution*. Columbia University Press, New York.