

# Return to Eden?

## Promises and Perils on the Road to a Global Superintelligence

Francis Heylighen  
*Global Brain Institute*  
*Vrije Universiteit Brussel*

**Abstract:** The concept of Singularity envisages a technology-driven explosion in intelligence. This paper argues that the resulting superhuman intelligence will not be centralized in a single AI system, but distributed across all people and artifacts, as connected via the Internet. This *global brain* will function to tackle all challenges confronting the “global superorganism”. Its capabilities will extend so far beyond our present abilities that they can perhaps best be conveyed as a pragmatic version of the “divine” attributes: *omniscience* (knowing everything needed to solve our problems), *omnipresence* (being available anywhere anytime), *omnipotence* (being able to provide any product or service in the most efficient way) and *omnibenevolence* (aiming at the greatest happiness for the greatest number). By extrapolating present trends, technologies and processes, the paper shows how these abilities may be realized within the next few decades. The resulting solution to all our individual and societal problems can be conceived as a return to “Eden”, the state of abundance and peace that supposedly existed before civilization. In such a utopian society, individuals would be supported and challenged by the global brain to maximally develop their abilities, and to continuously create new knowledge. However, side effects of technological innovation are likely to create serious disturbances on the road to this idyllic state. The most important dangers are cascading failures facilitated by hyperconnectivity, the spread of psychological parasites that make people lose touch with reality, the loss of human abilities caused by an unnatural, passive lifestyle, and a conservative backlash triggered by too rapid changes. Because of the non-linearity of the system, the precise impact of such disturbances cannot be predicted. However, a range of precautionary measures, including a “global immune system”, may pre-empt the greatest risks.

### The Singularity

It has been observed at least since the 1960s that technological advances seem to arrive ever more quickly. This forces upon us profound changes in our psychology, lifestyle and society (Heylighen, 2008; Toffler, 1970). A number of authors (e.g. Kurzweil, 2005; Vinge, 1993; Eden, Raker, Moor, & Steinhart, 2013) have argued that the advances are so fundamental that at some point in the near

future, they will culminate in a *technological singularity*. This concept of “Singularity” (in the short form) is used in two logically distinct, albeit linked, senses:

1. an acceleration of technological progress so radical that it appears like a discontinuity in what was up to now a continuous development;
2. the creation of an artificially intelligent computer system (AI) so intelligent that it can reprogram itself in order to become ever more intelligent, thus radically surpassing any human abilities.

The link is that extreme acceleration of technological progress makes the creation of superhuman intelligence more likely, while the creation of superhuman intelligence would radically accelerate further technological progress, given that this intelligence would invent new solutions much more quickly than any humans could.

Sense (1) is the original application—first suggested by John von Neumann—of the mathematical concept of singularity to the curve of increasing technological capabilities (Kurzweil, 2005). As these abilities increase more quickly than exponentially, their curve may take the shape of a hyperbolic growth, i.e. a function that reaches an infinite value after a finite time. At this point in time, progress would happen so quickly as to appear unlimited. Therefore, we cannot extrapolate the curve beyond this singular point. This makes us incapable of conceiving what will happen after this world-shattering event.

Sense (2) is the interpretation by the mathematician/science fiction author Vernor Vinge (1993) of this event as the emergence of a self-improving AI, initiating an explosion in (computer) intelligence. Both interpretations are supported by Moore’s law, which implies an exponential growth (but not a hyperbolic one) in the computing power of chips (Kurzweil, 2005). The assumption is that the technology-driven explosion in computing power would result in a concomitant explosion in information-processing ability, and therefore in intelligence. This seems to be the standard interpretation among Singularity theorists at present (Chalmers, 2010; Eden et al., 2013; Kurzweil, 2005).

### Why an autonomous AI Singularity cannot happen

I have argued earlier (Heylighen, 2012a) that the scenario of an explosion in AI through self-amplification is very implausible. The reason is that intelligence requires more than computing power: according to the *situated and embodied* theory of cognition (Clark, 1998; Steels & Brooks, 1995), intelligence must be embedded into the outside world. The higher the intelligence, the more sensitive and pervasive the “sensory organs” and “muscles” it would need in order to use its capabilities effectively. Such an immense array of finely tuned and coordinated sensors and effectors actively coupled to the environment is much more difficult to build than a faster chip.

A stand-alone AI computer, however powerful its processor and however clever its programming, would be little more than what embodied cognition theorists have called “a brain in a vat”: an entity that may be capable of complicated abstract reasoning, but that is not in touch with the real world (Heylighen, 2012a). The idea that such a system could autonomously boost its intelligence, in the sense of its power of making smart decisions about the world, is akin to the idea

that you could build a *perpetuum mobile*, i.e. a machine that would generate its own energy out of nothing (Ashby, 1981). Just as the law of conservation of energy prohibits a *perpetuum mobile*, Ashby's law of requisite variety (Ashby, 1958; Heylighen & Joslyn, 2003) prohibits a machine that could generate more specific information about the world than what it receives through its data feed. At best, it could generate more abstract models of the data, thus becoming able to predict some of its recurrent patterns. But it could not "invent" information about independent events that it did not have access to. The "no free lunch theorem" (Wolpert & Macready, 1997) formulates a similar restriction: in problem solving there are no universal shortcuts; you need knowledge about each specific situation. Yet another formulation of this constraint is the "knowledge acquisition bottleneck" (Cullen & Bryman, 1988): the key limiting factor in AI development is the speed with which knowledge about the world can be entered into the system.

In conclusion, an explosion in intelligence requires an accompanying explosion in the amount of information being acquired about the world, and therefore in the bandwidth of the channel through which the intelligent system interacts with that world. This makes a self-amplifying AI explosion very unlikely, because an AI program cannot extend its physical grasp of the world as easily as it could reprogram its informational routines. Unlike software development, building sensors and effectors is limited by a wide range of physical constraints. The necessary tools and resources can only be provided by a large group of people willing to invest a vast, coordinated effort into adapting and integrating these disparate physical systems into a complex and abstract software system—a development that is likely to take several decades. Only after this integration with its physical "body" would the AI be able to autonomously boost its knowledge of the world by starting to exploit its extended interaction abilities. Thus, the development of a superintelligent AI would be neither explosive, nor self-controlled.

An even more fundamental problem is that AI systems intrinsically have little sense of what is important, valuable, or worth doing: the problems they are supposed to solve are normally formulated by the user or programmer, not by the system itself. Therefore, such systems cannot autonomously decide and act outside of their narrow domain of programming. At best, robotic agents may be programmed to attain a system of interdependent goals (Steels & Brooks, 1995), such as mowing a lawn while avoiding obstacles and making sure their batteries do not run out. However, such preprogrammed goal hierarchies lack the endless nuances, the flexibility and the real-world significance that characterize our own value systems. These values are embodied in our genes, neurons, synapses, hormones, neurotransmitters and other components of our physiology, which together control our sensitivities, emotions, attitudes and behaviors. As such, they are the product of billions of years of biological evolution, millennia of socio-cultural development, and decades of personal experience (Heylighen, 2012a).

The fact that our goals and values are much more complex than we think is shown by plenty of observations from ethics, psychology and neurophysiology (e.g. LessWrong, 2013; Muehlhauser & Helm, 2012). In a typical fable illustrating this problem, a genie grants an apparently lucky guy three wishes. Formulating desired outcomes (i.e. goals) seems simple enough: perhaps everything he touches should turn to gold! The literal-minded genie then produces the formulated condition. However, it quickly becomes clear that this change also affects other unstated conditions in a way

that is more negative than positive (like starving because you cannot eat gold). The story ends with the guy scrabbling to formulate a third wish that would undo all the damage produced by the fulfillment of his first two wishes, which he had formulated too simply, neglecting to take into account their effects on unstated values (like the value of being able to eat).

This problem cannot be resolved by just thinking a little longer and deeper about what it is precisely that you want: most of the conditions that you would rather have or avoid are subconscious, and cannot be captured in a simple formula. In fact, formulating a realistically complex system of goals is a problem beyond the limits of physical computability (Ashby, 1972; reprinted in: Ashby, 1981): the corresponding “utility function” (mathematical measure of value) simply has too many degrees of freedom.

The Singularity theorists (e.g. Muehlhauser & Helm, 2012) who are aware of this problem see it as an obstacle to the design of a “Friendly AI”, i.e. a superintelligent AI agent that could be depended upon to only make decisions beneficial to humans. Like in the case of the genie, a poor formulation of the goals programmed into the AI could lead to all kinds of catastrophic side effects, thus making the AI effectively “unfriendly”. But the more basic problem is that an AI programmed with a too simple value system would not be able to deal with the complexities of the real world, because it would not know how to recognize significant opportunities, challenges or dangers, or how to choose wisely among the myriad of options that it would have in any concrete situation. And since there is no way to develop such a utility function by mere computation, such an AI would be essentially helpless without human guidance when taken outside of its formal domain of specialization...

### A distributed Singularity

While an AI system cannot bootstrap itself into superintelligence without an impracticably large amount of human help, we can foresee the self-amplifying explosion of a different kind of intelligence. Its requisite interaction channel already exists, in the sensory organs and muscles of billions of people, supported by billions of technological systems, distributed across the planet. Its utility function is implicit in the trillions of value judgments these people make each time they choose between different things to do, to acquire, or to pay attention to.

What is needed to make this immense sensing, valuing, and acting capability available for intelligence is to *interconnect* these people and tools via a vast and powerful network. That network too already exists, in the form of the *Internet*: any information collected or decision made by any human or technological agent(s) can in principle be transmitted nearly instantaneously to any other agent(s) for further processing and/or enacting. The amount of information propagating across this global network is many orders of magnitude larger than the most powerful channel we could envisage to feed a single AI, however advanced its processing capacity.

Just like the neural networks in our brain, the global network processes this information in a *distributed* manner, with billions of human and technological “neurons” working in parallel—on partly the same, partly different data—while aggregating their results into collective decisions and

actions. Unlike the centralized, tightly controlled systems imagined by AI theorists, this network functions like a *complex adaptive system* (J. H. Miller & Page, 2007): a decentralized, self-organizing whole formed by an immense variety of interacting agents. This makes the overall system much more powerful, robust, and adaptive than a preprogrammed central processor that takes in information sequentially. As it adapts, this network learns from its experience by expanding and extending the most useful technologies, ideas, links, institutions, and agents, while setting aside and eventually forgetting the less useful ones. Moreover, it learns to ever better coordinate the activities of its components, by creating ever more direct and more powerful links between them. It thus increases its own intelligence. This increase can only accelerate as a more powerful and intelligent network also expands and develops more quickly. Therefore, the scenario I advocate proposes an explosion in *distributed intelligence* (Heylighen, 2012b).

In this scenario, data gathering, processing and valuing are spread across the planetary network of people, computers, and other technologies. This system of collective intelligence has been called the *Global Brain* (Bernstein, Klein, & Malone, 2012; Goertzel, 2002; Heylighen, 2008; Mayer-Kress & Barczys, 1995). It emerges through a self-organizing coordination between the different people, artefacts and links in the network. This coordination accelerates as communication channels, shared memories, interfaces, and our general understanding of distributed intelligence advance—a phenomenon illustrated by the astonishing growth in the size, capabilities and impact of the Internet over the past decades. General-purpose AI, in contrast, has witnessed only modest progress over these same decades—in spite of the perpetually bold forecasts and the unrelenting growth in processing power expressed by Moore’s law.

### The Singularity as a metasystem transition

While advances in AI systems will undoubtedly contribute to accelerating technological progress, for me it is clear that the bulk of innovation will come from the network-mediated interactions between people and their technological supports. Given the physical and biological nature of these components, which are still subjected to the kind of inertia and friction that has been practically eliminated in computing, I do not think that the speed of advance will ever approach infinity. Yet, the concept of Singularity remains useful as a metaphor to help us understand the absolutely radical nature of the changes to be expected.

The emergence of the Global Brain (GB) can be seen as a transition to a higher level of complexity or evolution (Heylighen, 2000; Maynard Smith & Szathmáry, 1997), a process that Turchin (1995) has called a *metasystem transition*. Such a transition separates two qualitatively different phases of organization, like the transition from gas to liquid, from chemical to biological organization, or from single cell to multicellular organism. That is why it is very difficult for us to look beyond the Singularity and try to imagine what the new organization will be like.

While such a transition may appear like a discontinuous, singular event when seen from afar, zooming in reveals that what appeared like a step function is better approximated by a smooth, sigmoid (S-shaped) curve. Such a growth curve starts by increasing exponentially until it reaches a maximum speed, after which it slows down and stabilizes at a much higher level. This represents the

natural dynamics for any spreading innovation—whether in technology or in biology (Modis, 2012). Therefore, it is likely that the present transition too will be characterized by a maximum speed, which will depend on the inertia and friction of processes in the real world. Arguably, we are already near that maximum speed stage (Heylighen, 2012a). This would imply that while the speed of change will continue to be very high for the next few decades, it is on the verge of a long slowdown towards a new, more stable regime.

Like the other chapters in this volume, the present paper focuses on the developments and trends in the coming period leading up the Singularity, with an emphasis on their benefits and dangers to society. However, I will also propose a first glimpse at what lays beyond the transition, and try to imagine what a Global Brain regime may be like.

### **Abilities of the Global Brain**

According to the theories of cybernetics (Ashby, 1964; Heylighen & Joslyn, 2003) and of situated and embodied cognition (Clark, 1998; Steels & Brooks, 1995), the function of intelligence is not abstract reasoning, thinking, or computing. It is rather directing and coordinating the actions of an organism within its environment. All organisms have evolved to survive and grow, by evading dangers and exploiting opportunities. This process can be summarized as “tackling challenges”, where a *challenge* is any situation that threatens with a loss of fitness (danger) or promises a gain in fitness (opportunity) (Heylighen, 2012b, 2012c). Thus, *a challenge invites an agent to act*, in order to realize the gain and/or avoid the loss. The intelligence lies in the conception and selection of the most effective combination of actions to execute for any given situation.

Intelligence, in this perspective, is the ability:

- (1) to recognize (*perceive*), interpret (*process*) and prioritize (*value*) meaningful challenges;
- (2) to conceive, select, and initiate the right actions for dealing with them.

“Meaningful” here means relevant to *fitness*, i.e. the long-term ability to survive, develop and grow within the organism’s complex and variable environment. The highly multidimensional function of fitness is the ultimate value for any system that desires to survive and thrive. However, this value function is not a priori given—unlike the utility functions used to program AI agents. It has to be learned by the organism through myriad processes of trial-and-error across evolutionary time. Our individual and social values are the accumulated results of these on-going learning processes.

The global brain perspective focuses on intelligence and fitness at the level of planetary society: the whole of humanity together with its artifacts and the ecosystems that it controls and relies on. This system of mutually dependent people, technologies and ecosystems defines the *global superorganism* (de Rosnay, 2000; Heylighen, 2007a; Stock, 1993). A superorganism is an organism with components that are organisms themselves. However, note that organisms typically extend their reach to include non-organic components as well (Dawkins, 1999; Turner, 2000). For example, tiny polyps produce coral reefs for their housing and support, while beavers create dams and artificial lakes to protect their nests. Similarly, humans extend both their bodies and their minds (Clark &

Chalmers, 1998) by means of technological artifacts, and this at both the individual and the collective level.

According to living systems theory (J. G. Miller, 1995), the social systems they form in this way exhibit the same basic functions of life as their individual organisms: ingestion, processing, distribution, excretion, storage, etc. This applies in particular to global society (Heylighen, 2007a), which is the only social system that has a true external boundary—space.

The Internet plays the role of the nervous system for this superorganism. The accelerating development of this network and its supporting information processing technologies is turning this nervous system into a brain-like intelligence that can deal with increasingly difficult challenges (Heylighen, 2007a; Heylighen & Bollen, 1996). Thus, the thrust of change on the brink of the Singularity is towards increasing distributed intelligence. This will augment our ability to deal with any kind of meaningful challenge, individual as well as societal. Given the immense, and accelerating, reach and power of a globally distributed intelligence, it is clear that its capability for solving problems will far surpass anything we could have imagined so far.

But what exactly will those capabilities be? In the end, something can be expected to evolve when there is a *selective pressure* to produce it. The concept of “selective pressure” is a shorthand for the following mechanism. Assume many variations of a given design (e.g. an organism or an artifact) that are not all performing equally well. If we look closer, we may note that the ones doing best satisfy some criterion that makes them more likely to survive and proliferate in the given environment. For example, rabbits that can run faster are less likely to be caught by foxes. Thus, natural selection tends to keep the faster rabbits, while eliminating the slower ones. As variation and selection carry on, subsequent generations of rabbits tend to become faster and faster. It is as if there is a “pressure” pushing the rabbits towards higher speeds. That observation allows us to predict that future generations of rabbits will become ever faster—up to the point where further increases are no longer possible given the physical limitations of rabbit physiology.

The same principle helps us to forecast the evolution of the information and communication technologies (ICT) that enable the global brain, and the manner in which they will be used. Researchers are constantly devising new tools, methods and technologies, or new variations on existing technologies. The selective pressure here is that technologies tend to be adopted when they help people solve their problems—no matter what these problems are. Variants that satisfy this criterion better will maintain and spread. Variants that satisfy it not so well will eventually be discarded and forgotten. Thus, technological innovations are competing in a race to become ever more helpful to society.

Let us subdivide this general criterion of “helpfulness” into more specific attributes. A technology will be more helpful if it is:

1. more intelligent and knowledgeable, so that it can find solutions to more complex and diverse problems
2. more widely accessible, so that you can use it whenever you need it
3. more powerful and efficient in realizing its solutions
4. more beneficial in the impact these solutions have on individuals and society

These criteria determine distinct selective pressures for the evolution of ICT and of the institutions that interconnect them with each other and with their users—i.e. the global brain. The acceleration inherent in the Singularity implies that this evolution will quickly reach a level beyond anything that presently exists. Since the processing and distribution of information does not have any obvious physical limitations—unlike rabbit physiology—there does not seem to be anything to stop this race towards ever-higher capabilities. Let us therefore consider the limit of *infinite capability*. The “helpfulness” attributes would then turn into the “divine” attributes that characterize the God of monotheism (Mann, 1975):

1. *omniscience*: being able to answer any question or solve any problem
2. *omnipresence*: being accessible everywhere at every moment,
3. *omnipotence*: being able to produce any effect or achieve any goal,
4. *omnibenevolence* (or *perfect goodness*): being willing to help (and not harm) everybody.

It is not surprising therefore that both religious (Teilhard de Chardin, 1959) and non-religious thinkers (Otlet, 1935) have compared the emerging global brain to a God in the process of formation (Heylighen, 2011).

On the other hand, we know from logical and scientific limitation principles—such as the theorem of Gödel or the law of energy conservation—that infinite capabilities are intrinsically impossible (Barrow, 1998; Yanofsky, 2013). For instance, a classic paradox of omnipotence (Hoffman & Rosenkrantz, 2012) is the question whether God can create a rock so heavy that He Himself cannot lift it: if He cannot, He is not omnipotent; if He can, He is not omnipotent either. Similar paradoxes beset the notions of omniscience (Wierenga, 2012) and omnibenevolence. The conclusion is that—like the rabbit—the global brain (GB) must remain limited in its abilities.

However, these limits are likely to be so far beyond our present limitations that we may as well see them as belonging to a wholly new level of reality. In that sense, the classic list of divine attributes can help us to imagine what such a “god-like”, superhuman intelligence would be capable of. However, in contrast to the way most believers see their God, we should be careful not to conceive of the GB as a separate, personal being, with its own will independent from the people that constitute it. The GB is not a “higher power” that we should defer to. It is merely a particularly effective network of self-organizing interactions supported by ICT, in which we all play our part.

Thus, while religious analogies may be helpful for conveying the sheer magnitude and impact of these developments, I have no intention whatsoever to suggest that theology or Scripture can offer a guideline for understanding our future, as some Singularity thinkers seem inclined to do (Proudfoot, 2012). On the contrary, I hope to show that a critical and pragmatic, but open-minded, attitude is our best hope for coming to terms with a transition that is bound to change our worldview, culture and society to their very core. Let us, first, try to conceive the promises of these changes, then, consider the accompanying dangers and negative side effects.

## The global brain's promises

The future applications of the distributed intelligence technologies underlying the emergence of a GB are virtually infinite in their variety. Hence, we cannot conceive them with any degree of accuracy or comprehensiveness (Heylighen, 2013). Let us then use a pragmatic interpretation of the list of “selective pressures” or “divine attributes” as a simple classification scheme for investigating some of the more “Singular” developments to be expected.

### Omniscience

The most obvious attribute to expect from a post-Singularity, superhuman intelligence is a *practical omniscience*: the knowledge of everything needed to help humanity deal with its challenges. At this moment, the web already gives access to about every piece of knowledge that was ever published. However, this knowledge is quite fragmented, and it is often difficult to separate the wheat from the chaff, or to get a reliable answer to a concrete question.

The most successful attempt to date at organizing this knowledge is *Wikipedia*, the global encyclopedia that is being read, written and edited by millions of volunteers worldwide. The basis of its success is *stigmergy*, a mechanism of self-organizing coordination between independent agents (Heylighen, 2008; Parunak, 2006). The principle is that the work of one individual (e.g. an edit of a Wikipedia text) leaves a public trace (e.g. a change in the corresponding web page) that can stimulate one or more other individuals to continue the work (e.g. add further details or correct wrong assumptions). Thus, independent contributions build further the one on top of the other, producing a collective result much richer and more complex than could have been achieved via any traditional, centralized form of organization. There are no obvious limits to such a self-amplifying process of contributions eliciting further contributions. Therefore, we should expect that Wikipedia and related initiatives will continue to expand and deepen over the next decades, until they offer practically the whole of human knowledge in a clear and elegant presentation.

Providing knowledge in the form of an encyclopedia is still but a step towards omniscience: this knowledge should not just be available for reading and interpretation, but for directly solving problems and answering questions. Another important step is *structuring the knowledge* in such a form that it can be used for automated reasoning, without the need for humans to interpret typically ambiguous texts. This is the objective of the *Semantic Web* (Berners-Lee & Fischetti, 1999), a set of protocols for representing knowledge in a precise and dependable way, so that it can be interpreted and applied by machines.

While the Semantic Web initiative is older than Wikipedia, its advance has been much slower. The reason is that it is much more difficult than one would naively expect to formalize our largely intuitive knowledge. The fundamental problem is that the world is not composed of the kind of stable, objective categories that most scientific theories and AI “ontologies” presuppose: real-life phenomena are intrinsically variable, fuzzy and context-dependent (Gershenson & Heylighen, 2005; Heylighen, 1999). This implies that the semantic web will continue to move ahead relatively slowly, and that it will never succeed in capturing all of human knowledge in an explicit and logical format. However, its advances will eventually enable a greatly automated system of inference, allowing

people to get answers to questions (e.g. “are penguins warm-blooded?”, “which US presidents were born in a year divisible by 7?”, “which car mechanics specialized in Mercedes and speaking Russian work in my city?”) that are not readily available as such in Wikipedia or other databases, but that can be logically deduced from these data.

The rigidity of semantic web-like technologies can be complemented by more adaptive technologies, such as neural networks, machine learning, and recommender systems. Their underlying algorithms are able to give approximate answers to ambiguous questions, by aggregating a huge amount of fragmentary, subjective and imprecise data, while extracting the most important underlying trends and associations. They moreover can take into account the context in which the question was asked. For example, the same question asked by different people at different moments may require different answers. By considering the characteristics and earlier activities of the inquiring person, the system can guess which answers are most appropriate in the situation (Heylighen & Bollen, 2002). Moreover, such system can provide “solutions” (or rather suggestions) without any question being explicitly asked—again by taking into account the context of the activity.

Such systems have proven their usefulness in a variety of settings, e.g. when prioritizing results to a Google query, suggesting YouTube videos, or recommending books to Amazon customers. However, the underlying algorithms tend to be obscure, disjoint and hidden behind trade secrets. Over the next decades, we can anticipate that flexible and robust solutions will become available as open-source software and that standards will emerge so that they can be universally applied to guide any query. At that stage, every person using the Internet will be constantly guided towards the best solutions to that person’s problems, in most cases even without having to ask any questions (Heylighen, 2007a; Heylighen & Bollen, 2002).

The most frequently needed knowledge should not just be available via the network, but as much as possible inside people’s own brains, so that they can immediately apply it to the situation at hand. This can be achieved by *learning*. Web technologies are starting to make education much easier, more effective and more enjoyable (Heylighen, Kostov, & Kiemen, 2013). They make it possible to accurately tailor teaching to the learner, by providing the right challenges at the right moments, depending on the interests and abilities of the individual. Moreover, they provide immediate feedback on how well the learner is doing, thus stimulating further advances and catching problems at the earliest stage. Finally, they can encourage collective learning, in which students of different levels help each other with questions, suggestions and feedback. This is already achieved to some degree by “Massively Open Online Courses” (MOOCs, Rodriguez, 2012), such as the Khan Academy, Coursera, and edX, and by an endless variety of apps that use gaming techniques to make learning fun (Michael & Chen, 2005; Thompson, 2011).

At this moment, an enormous amount of energy is invested by volunteers, communities, schools, universities, companies, and governments in order to develop attractive and effective course material on about any possible subject. For the most important topics, material will be free to use by anyone in the world. Learners will be tested on their degree of advance in a relaxed, game-like, interactive manner, with the option to eventually convert their earned “points” into formal degrees. The only thing needed to fully exploit these possibilities is motivation from the learner. Next to people’s natural curiosity and desire to advance in life, their motivation will be boosted by powerful

mobilization techniques, such as gamification (Heylighen et al., 2013), and by the fact that if everyone around you has a solid education, then you are left in a very awkward position if you do not catch up...

A straightforward extrapolation of these developments tells us that in the near future, people anywhere in the world will be studying and obtaining degrees on any subject by following courses across the net. This will elevate the education level of humanity to such a degree that it is as yet difficult to ascertain the consequences: imagine a world where every minimally gifted adult has the equivalent of a PhD degree in an advanced domain, while having a broad and deep general education covering a wide variety of topics... At the very least, this will remove the sheer ignorance; lack of skills; deference to authority, convention and tradition; passivity; and lack of self-confidence that has prevented so many people from developing their innate capabilities.

Furthermore, it is likely to produce an explosion in *creativity*, as billions of highly educated people collaborating closely with the “omniscient” knowledge bases and sensor networks of the GB and each other will explore an unimaginable range of issues, ideas and solutions. This creativity will be further amplified by algorithms for knowledge discovery, machine learning, and data mining, which search for general patterns in immense arrays of concrete data. They are likely to be enhanced by novel theoretical insights into the nature of knowledge, complexity and creativity. The synergy between the automated collection of data, the computational search for regularities, and the hunches and intuitions of highly skilled people will multiply our capabilities for research. This will turn the scientific discovery of laws, concepts, and theories into a routine activity—as straightforward for the GB as writing messages is for us now. Thus, the GB will be able to create new knowledge on the spot—whenever it may be needed to address a particular challenge or question.

An essential part of that knowledge will be about the superorganism as a whole: how does the world function, and what opportunities and dangers can we anticipate? This requires, first, a general, transdisciplinary theory of complex systems and their networks of interactions: a *global systems science* (Helbing, Bishop, Conte, Lukowicz, & McCarthy, 2012). This theory would then be implemented as a *living earth simulator* (Helbing et al., 2012): a comprehensive computer model of global society, the planetary ecosystem, and the way they interact. Such a simulation would allow us to anticipate the effects of different events or interventions on the global superorganism. That would make it possible to formulate a nearly optimal trajectory for global development, while as much as possible steering clear of conceivable obstacles, catastrophes, and negative side effects (Helbing, 2013). However, in order to remain up to date and to adapt to unforeseen disturbances, this model would need to be constantly fed with new data. Moreover, it should be able to test out its hypotheses in the real world, so as to correct possibly inaccurate assumptions. As we noted earlier, this would require a very high-bandwidth interaction network, which (Helbing et al., 2012) call a *planetary nervous system*. That brings us to the next “divine” attribute.

## Omnipresence

A global brain should be able to monitor what is happening everywhere in the world at any moment, so as to extract the relevant information for anticipating what will happen next. Moreover, it should be able to intervene anywhere anytime, in order to tackle challenges the moment they appear.

This capability already exists to some degree thanks to the nearly universal availability of mobile phones. These allow people anywhere in the world to report on what they are observing, or to receive instructions on what to do. Soon, all these phones will be replaced by *smartphones*, with in-built sensors that can send and receive more fine-grained data (e.g. documents, photos, videos, GPS coordinates) about what is going on at their locations. This information can be entered into a global, shared database that keeps track of what is happening where and when. Such a global picture of the situation makes it easy to coordinate actions and dispatch help to the places most in need, e.g. in situations of disaster relief (Gao, Barbier, & Goolsby, 2011).

The most advanced interfaces between person and GB do not even require the manipulation of a smartphone: they can be embedded in people's clothes (*wearable computers*, Barfield & Caudell, 2001) or glasses (e.g. "Google Glass," 2013). Thus, they are "always on", providing context-sensitive information together with the situation as perceived by the user—wherever that user is. The result has been called *augmented reality* (Van Krevelen & Poelman, 2010), as the extra information provided by the GB is superimposed on the person's experience of the real world. The final stage will be a *brain-computer interface* (He, Gao, Yuan, & Wolpaw, 2013; Zander, Kothe, Jatzev, & Gaertner, 2010), in which sensors directly read people's intentions from their brain waves or other physically measurable signals, while responding by generating similar neural signals. At that point, people will merely need to think about something to see the GB answer their questions or realize their desires.

A different aspect of omnipresence is the emerging *Internet of Things* (Atzori, Iera, & Morabito, 2010). This refers to the addition of simple wireless sensors (e.g. RFID tags) to various physical objects, together with a communication protocol for consulting these sensors across the Internet. This makes it possible to monitor the state (e.g. location, temperature...) of each object remotely. Thus, the global brain's awareness expands from people to things. Eventually, thanks to inbuilt processor chips, these objects will be able not only to passively send out information, but to intelligently communicate with people and with the network. This development has been called *ubiquitous computing* or—perhaps more accurately—*ambient intelligence* (Bohn, Coroamă, Langheinrich, Mattern, & Rohs, 2005). It refers to the vision of a physical environment that intelligently responds to the needs of the people present. For example, when you walk through a building at night, lights switch on automatically in the spaces ahead, doors open and music begins to play, while these activities switch off behind your back.

For adequate response, objects should not just have sensors to receive data (e.g. a movement detector), but actuators or effectors that convert instructions into physical action (e.g. a door opening mechanism). The more autonomous of such devices may be called "robotic". However, they are unlikely to take the humanoid shape that most people still associate with robots. More typical examples are the drones used by the military: small, remotely controlled flying vehicles that can

monitor vast 3-dimensional spaces (e.g. in search for a hiker lost in the mountains), while being able to intervene in some way (e.g. providing the hiker with water and a first-aid kit).

Many such tiny robots can form a self-organizing “swarm”, wirelessly passing on information to their neighbors (and eventually to the GB) in order to coordinate their activities. Such “mesh networks” are extremely robust, and would remain functional even in situations where many of the devices have stopped responding because of damage or poor communication (Akyildiz, Wang, & Wang, 2005; Dressler, 2008). This would ensure that the GB would be accessible even in the most difficult circumstances (such as an explosion in an underground tunnel)—i.e. wherever its presence might be needed.

Giant swarms of miniature sensors may be distributed ubiquitously across the cities, lands, oceans, ice sheets, and even atmosphere of the Earth, e.g. by “spraying” them from robotic planes. They would be so small that they could extract the tiny amount of energy they need directly from their environment (e.g. from sunlight or waves), thus functioning autonomously. By using robust peer-to-peer protocols for wireless communication, they could propagate the relatively few data each sensor gathers via their neighbors’ neighbors until they reach the GB network, which would aggregate the bits and pieces into a coherent picture. This would allow the GB to monitor the planet as a whole, not just, as is presently done, from the perspective of a far-away satellite, but from the inside out. If a problem is detected (e.g. a hot spot near a dry bush, or the presence of toxins), the GB would immediately dispatch the nearest robotic vehicle or person to deal with it (e.g. douse the spot with water so as to prevent fire).

## Omnipotence

The omnipotence attribute may seem the one least fitting to describe the GB, given that it implies acting in the physical world rather than in the virtual world of information. Therefore, it is subjected to the law of conservation of matter and energy, which states that you cannot create physical resources out of nothing. While this limitation is insurmountable in principle, in practice distributed intelligence can dramatically increase the availability of resources. It achieves this by *mobilizing* them—that is, directing them in a precise, coordinated manner at the appropriate targets (Heylighen et al., 2013). At present, most physical resources—such as food, water and energy—are consumed by dissipation and waste, not by productive use. Technological progress has spectacularly diminished such losses (Heylighen, 2008). Yet, waste remains far too common. A true GB should be able to achieve a maximum useful effect with a minimum of resources. This would eliminate any practical limitations on production and control.

The principle is that waste is ultimately the generation of *entropy*. This means that energy and resources are diffused or *dissipated* during the process, so that they are no longer available for productive work. The second law of thermodynamics states that every process is accompanied by some degree of dissipation, but puts no lower bounds on that degree. On the other hand, information is in a sense the negative of entropy (Heylighen & Joslyn, 2003): it reduces the uncertainty or disorder that characterizes entropy by determining the precise state of the system. Therefore,

providing the right information can minimize entropic diffusion, by specifying the optimal course for the process at every moment. In theory, an omniscient system should be able to steer the process along a path with virtually no dissipation, by directing its course in every microscopic detail.

A practical illustration of this principle can be found in *3D printers* (Lipson & Kurman, 2013). These printers can very precisely control the position of a print head in space, and the amount and type of material (“3D ink”) ejected from it at every moment. Their capacity for minute adjustments will only continue to increase as we move into *nanotechnology*—the manipulation of matter at the scale of cells and molecules (Drexler, 2013). In principle, such a printer (or nano-assembler) can produce any material tool or object, however complex, if it receives the right instructions about where in space to deposit which tiny bit of which material. These instructions specify the optimal path by which the different bits are to be assembled, layer-by-layer, into the desired shape.

Compare this to the traditional methods of construction. Here, large chunks of various materials are melted or cut into pieces, after which all the superfluous bits of material are shaven off, until each piece has the right shape. Then these pieces are assembled via a complex process of fitting, screwing, and welding. All of these operations are accompanied by a great waste of energy and materials—unlike the precise printer movements that add just what is needed where it is needed.

This meticulous control is possible because of the detailed information in the instructions that direct the print head. These instructions are generated by the printer’s processor from a 3-dimensional specification of the required object. Such an informational blueprint will typically be downloaded from the Internet, where it was developed and published by a community of designers (Anderson, 2012). These designers will rely on typical GB mechanisms: Internet-supported collaboration, advanced software, libraries of common shapes and specifications, collective experience, and prototype printouts to test out in the real world so as to refine the design. Whoever needs a specific object can then have it physically produced it within minutes by following a simple procedure:

1. select an appropriate design, possibly with the help of the GB’s recommendations;
2. if necessary, adapt it to specific tastes or requirements (e.g. color, size, material) by minor edits of the file;
3. send it to the printer;
4. collect the “printout”.

In a decade or two, high-quality 3D printers will be as ubiquitous as computers or TVs are now (Anderson, 2012). At that stage, anyone anywhere will be able to immediately get about any not too complicated object—for a cost little more than the one of the required 3D “ink”. Thus, the development and distribution of 3D print designs via the Internet will eliminate all the most wasteful processes in traditional industrial production:

1. have a team of in-house specialists research and develop the production process;
2. collect the materials and components from different suppliers;
3. assemble these components in large-scale, energy-intensive and labor-intensive factories;
4. market the products to potential clients;

5. ship and distribute them worldwide via warehouses, retailers and other middlemen to the eventual buyers.

Not just printers and other assemblers can be directed remotely via distributed intelligence: people, robots, and any Internet-enabled objects that have effectors can be mobilized to perform the right actions at the right time in the right place, and this in a smooth, coordinated manner. The many communities of volunteers on the Internet illustrate that it suffices to formulate a worthwhile objective and to provide an effective collaboration platform to get people to work efficiently towards the most ambitious targets (Heylighen, 2007b). These include writing down all the world's knowledge (Wikipedia), programming the software to run the Web (Apache, Linux, ...), or supporting start-ups, pro-democracy movements, or charities with money, ideas and volunteer work. An extension of the underlying techniques for encouraging and coordinating such beneficial activity (Heylighen et al., 2013) should allow the GB to harness people worldwide to tackle any kind of challenge.

Moreover, the GB will be able to remotely control the most important robotic devices, such as vehicles, machines, building equipment, cameras, etc. The combination of human and robotic power will allow it to tackle large-scale, distributed problems, such as disasters (Gao et al., 2011), famines, or epidemics, in the most efficient manner. On a more everyday level, the GB will regulate all of the world's traffic, including cars driven by people or by robotic systems, public transport, planes, and the logistic flows of goods and materials across the world via trucks, trains and ships. This would minimize all waste of time and energy caused by needless journeys, detours, traffic jams, or less than optimally loaded vehicles. Similar gains can be expected by the precise direction and coordination of agricultural and industrial production, so as to minimize pollution and resource consumption and maximize productivity.

Next to its physical dimension, omnipotence has a cognitive dimension: the ability to solve any practically conceivable problem by planning an effective course of action for dealing with it. This capability is already inherent in the omniscience and omnipresence attributes: the GB would have all the knowledge it needs and be able to test any suggested actions immediately in the real world to check whether they are effective. Coupled to an intelligence, creativity and computational power orders of magnitude larger than what we have now, this should be sufficient to come up with strategies that can effectively tackle the most difficult challenges. These strategies can then be directly implemented via the network of agents, sensors and effectors that the GB commands.

The conclusion is that a GB-regulated society would be one of *abundance* (Diamandis & Kotler, 2012; Drexler, 2013). Any remaining shortage or scarcity would have vanished, as the GB can provide any desired information, product or service anywhere, at a negligible cost in labor, energy or raw materials. Moreover, through effective, coordinated action, such a society would be able to tackle all its significant challenges, both on the societal and on the individual level. The only question that remains is whether the GB would effectively *want* to tackle all these problems...

## Omnibenevolence

The last “divine” attribute is perhaps the least evident one. *Omnibenevolence* means that the GB would be universally well intentioned, ready to help any individuals or groups achieve benefit, while minimizing any harm they may experience. The problem is that helping one individual or group may harm another one, as they have conflicting interests. This can be illustrated by various ethical dilemmas, where the issue is raised whether you could or should sacrifice N individuals if this would be necessary to save M different individuals.

Perhaps the most general solution is the one proposed by *utilitarian ethics*: imagine that there exists a function that can calculate the total amount of benefit (positive value) minus the harm (negative value) for humanity as a whole. Call the result of this calculation *global utility*. Then choose the action that maximizes global utility, no matter which harm it may inflict on certain people, because you know that any other action will inflict more harm, and/or produce less benefit. In this way, your ethical choices are guaranteed to produce *the greatest happiness for the greatest number of people*.

The problem with utilitarian ethics is that nobody knows exactly what that function is. We already noted there is no way to compute it (Ashby, 1972). The GB can to some degree tackle that problem by aggregating or inducing a collective utility function from all the explicit and implicit value judgments made by its billions of human components. However, it would be impossible to calculate the optimal value of this function by applying it to *all* the potential courses of action the GB could follow. You might expect that such a calculation would be straightforward for an omniscient GB. However, here we run into the classic limits on computability that make it impossible to explore all the options in an exponentially exploding search space—a space of possibilities which in this case would be absolutely staggering in its complexity.

In practice, therefore, such complex optimization problems are not solved by a deterministic calculation, but by means of a computer simulation that mimics distributed self-organization or evolution. First, the different possibilities are optimized locally, by choosing the apparently best action in each concrete case without taking into account its (potentially negative) effect on other actions. The different possibilities for actions are then randomly juggled through a myriad of different combinations, while trying to adjust the one to the other, until they settle into a nearly optimal configuration. In general, you cannot prove that this configuration is the one with the globally highest utility. But after trying out countless large and small variations on it, you may be confident that you are unlikely to find something much better. The GB will probably apply some more sophisticated variation of this method to select the actions that maximally benefit humanity as whole.

Still, many people fear that an omniscient computer system would enslave or eradicate humans rather than help them, so that it can take full control of the planet. Therefore, the last fundamental question is: why would the GB want to benefit humanity in such a way? The simple answer is: because humanity is part and parcel of the GB! The GB, by definition, is the nervous system of the global superorganism, which is itself constituted by people and their technologies. The human components of this system cannot be replaced by mere artifacts because they offer much

greater flexibility, better judgment, and more sophisticated sensors and effectors than robots (Heylighen, 2012a). Eliminating them would not only destroy the major part of the superorganism, but greatly reduce the bandwidth of the GB-world interaction, its store of values, knowledge and experience, and therefore its omnipotent and omniscient capabilities.

The more positive reason for the GB's intrinsic benevolence towards humans is that by benefiting them it benefits itself—and vice versa. Because it is for an essential part constituted by people, optimizing humans' physical, mental and social capabilities optimizes the GB's own capabilities. This is not just the outcome of a rational reflection about what the GB's objectives should be: the distributed intelligence network is evolving at this instant precisely because it is beneficial to those who use it. In the on-going competition between different methods, technologies, and institutions, the winners will be those that effectively produce the greatest happiness for the greatest number (Heylighen, 2013), because given a choice people will eventually replace a system that satisfies them less by one that satisfies them more. That is exactly what I meant by “selective pressures” pushing evolution in a predictable direction. Therefore, the trend towards universal benevolence and other “divine” attributes is built into the very evolutionary process we have been discussing.

While this may seem to contradict the traditional view of evolution as rooted in conflict and competition, a number of authors have recently shown that the long-term drive of evolution is towards cooperation and synergy (Corning, 2003; Heylighen, 2006; Stewart, 2000; Wright, 2001). The principle is simple. Two individuals desiring to consume the same limited resource are by default in competition: whatever the one gains, the other one loses. Such an interaction is called a *zero-sum* game or a *win-lose* situation. It typically ends in the survival of the one most fit to obtain the resource, and the elimination of the other. However, in many cases evolution will eventually stumble upon an arrangement that is *win-win* or *positive-sum*. The interaction then becomes cooperative or synergetic, as both parties can extract more benefit by working together than by acting independently. Such a situation will be preferred by natural selection, since even the strongest party will become fitter by reaping the benefits of collaboration than by spending its energy in fighting off competitors.

Thus, the blind-variation-and-natural-selection of evolution eventually increases cooperation (Stewart, 2000)—albeit very slowly. The emerging GB spectacularly accelerates this development. The reasons are the following:

1. positive-sum arrangements are typically more complex and more difficult to find than zero-sum ones. By boosting our ability to explore myriads of potential interactions with other people and artifacts around the world, and by helping us to select the most synergetic ones, the GB multiplies the probability of discovering highly beneficial win-win situations;
2. the default assumption that resources obey a zero-sum logic only applies to matter and energy, because these are subjected to a conservation law. It does not apply to information, which can be reproduced, and thus shared, without limitations. As the open-access philosophy reminds us, giving information to others helps rather than hinders you in the long term (Heylighen, 2007b). Information, and not matter, is the basic currency of the GB.

Because this resource intrinsically obeys a positive-sum logic, the GB has every reason to spread it as widely as possible, thus maximizing the sum of benefits for every party involved.

3. by moreover making physical resources—which obey a zero-sum logic—abundant, the GB eliminates the need for competition. When there is enough for everybody, any conflict over scarce resources vanishes.

This argument for the accelerated evolution of a universal “goodness” that would benefit everybody, while eliminating the fundamental sources of conflicts, may still seem too theoretical to convince you. Let us therefore look at the concrete facts of socio-economic evolution. Over the past few centuries, the technological revolution has brought about a spectacular improvement in human welfare. It is well established that people have become much richer, better educated, and longer living, and that this global development continues at a rapid pace (Klugman, 2010).

More unexpected, perhaps, is that people are also less likely to fall prey to what we consider as morally “evil”: murder, war, slavery, prejudice, suppression, dictatorship, and corruption. Perhaps surprisingly for those who follow the news (which typically focuses on the most spectacular instances of crime, terrorism, and war), the proportion of people dying through homicide or conflict, or living under non-free, undemocratic conditions has been spectacularly decreasing over the past centuries and decades (Heylighen & Bernheim, 2000; Mueller, 2009; Pinker, 2011). The underlying reason for this trend seems to be that people are much better educated and informed about everything going on in the world (Bernheim, 1999; Pinker, 2011). Therefore, they are less willing to tolerate the situations of abuse that before they were not aware of, or used to endure because they did not know a better way. Moreover, with greater wealth and greater interdependency, there is simply less to gain and more to lose by engaging in conflict or by exploiting others.

In sum, we see here the same dynamic of increasing communication leading to increasing understanding of self and others leading to increasing synergy that drives humanity and its global brain towards global cooperation (Stewart, 2000; Wright, 2001), universal morality (Heylighen & Bernheim, 2000; Pinker, 2011) and thus omnibenevolence.

## Paradise regained

Assuming the practical equivalents of omniscience, omnipresence, omnipotence and omnibenevolence, we should expect the GB to efficiently solve about all present global and individual problems and conflicts, while spectacularly enhancing our technological, social, cognitive and biological capabilities. To describe the resulting situation of peace, abundance and fulfillment, we may turn to another religious metaphor, *Paradise*, an idyllic state or place in which everybody is intrinsically happy.

However, the term “Paradise” is ambiguous, as it can refer either to the initial *Garden of Eden* where humanity originated, or to the final *Heaven*, the state of bliss to which people’s souls are supposed to migrate after their death. The metaphor of Heaven seems to have inspired many Singularity theorists (Proudfoot, 2012). They conceive it as some kind of Platonic realm of data and

computation, held in the memory of an immense, superhuman AI computer system. After their death, people’s brain patterns would be “uploaded”, or reconstructed, into this nearly infinite, abstract information space. Thus, they could continue to live forever, as disembodied, virtual personalities. This quixotic vision of computational resurrection and artificial bliss is sometimes described as “the rapture of the nerds” (Eden et al., 2013).

The utopian ideal inspired by the GB (Heylighen, 2002), on the other hand, is much closer to the *Garden of Eden*. The idea of a “Garden of Eden”, “Arcadia”, or “Golden Age”, which recurs in different religions and mythologies, refers to an original state of abundance and leisure before civilization. In this “state of nature”, people could live spontaneously, without having to toil, worry about the future, or feel constrained by socially imposed laws, rules and taboos. There seems to be a real substance to this myth: anthropologists have observed a similar lifestyle in the last remaining tribes of hunter-gatherers they were able to study (Charlton, 2002; Diamond, 1987; Gray, 2009; Sahlins, 2004). Our hunter-gatherer ancestors were apparently able to lead a relaxed, playful and spontaneous life in part because their population densities were small. Therefore, the resources they could find in their environment were in general abundant relative to their needs (Sahlins, 2004), and they had little reason for conflict or coercion.

This is no longer the case for the billions of people presently living on our planet. That is why we need technology in general, and an omnipotent GB in particular, to regain the original abundance by exploiting the available resources much more efficiently. In principle, the GB can eliminate all the problems of poverty, pollution, resource exhaustion, conflict, ignorance, superstition, and work-induced stress. Moreover, it should be able to restore our physical and psychological environment to a more natural state, without the need to abandon the comfort and security brought by technology. Thus, it would create a new Garden of Eden, offering an inspiring, relaxed and joyful life for all.

In such a society, people would no longer *have to work*, since practically all the essential jobs would be performed by GB-supervised machines. However, they would still *love to work*, because the tasks remaining for humans would be so intrinsically interesting and challenging that people would feel bored without them. That is because the GB would use people for the tasks for which they are biologically and psychologically best adapted. These include:

- nurturing and caring for people;
- looking after plants, animals and the natural world;
- manipulating irregular objects and materials;
- experiencing and interpreting complex, multisensory stimuli;
- inducing general wisdom from personal experiences;
- exploring and being creative;
- developing new concepts and inventions;
- conversing and collaborating with others.

Moreover, the GB would organize its interface with people in such a way that they would be maximally stimulated or motivated, and minimally stressed or confused by the work they are offered to do (Heylighen et al., 2013). The main incentive for good contributions will probably not be money, as material reward will lose its value in a society where everything is abundant, but public *recognition*, in the sense of constructive feedback and an increased reputation or status. Reputation

mechanisms are already being used very effectively to motivate people in open collaboration platforms on the Internet (De Alfaro, Kulshreshtha, Pye, & Adler, 2011; Mamykina, Manoim, Mittal, Hripcsak, & Hartmann, 2011).

Attaining such a society will obviously require revolutionary changes in our present political and economic system. At this stage, it is difficult to predict how such changes will come about—although many commentators would agree that the present political system is unsustainable, and that change is bound to happen. The highly developed social democracies of Northern Europe, Canada, and Australia may offer a (still imperfect) model for such a future system of governance. Social democracy combines strong personal and economic freedom with an important role for the government (a centralized system that will eventually have to be replaced by a more participatory system of distributed governance). The income generated from taxes (ideally on wasteful consumption and pollution) is used to support people and processes that tend to be neglected in a pure free market economy because they do not generate profit. Such redistribution reduces inequality and increases overall welfare.

For example, in this system teachers and academics are normally paid by the government. Thus, they can provide high-quality education and scientific research freely to the whole of society—instead of only to those rich enough to pay for them. In this arrangement, a tenured professor has the academic freedom to research any subject in any manner of her choosing, without having to worry about being paid. However, she remains motivated to get the best possible results because she wants to improve her scientific reputation—e.g. through highly cited publications. Such an arrangement can be seen as a model for the comfortable but highly stimulating work life under a Global Brain regime, where most people would have an education level that allows them to do autonomous research.

But social democracies also care for the less talented. People with physical or mental handicaps, psychological illnesses, or with an education so poor that they cannot find a job under normal market conditions, are often employed in “sheltered” or “social” work places. These are organized so that even the most poorly skilled people have the support they need to perform simple but useful jobs (e.g. ironing clothes, or preparing goods for recycling). Moreover, these people earn a decent wage thanks to government subsidies that otherwise would have gone to welfare payments.

In a sense, the post-Singularity society could be seen as one huge “sheltered environment”. Here, everyone would be doing meaningful work at his own level of skills, supported by the GB whenever needed, without having to worry about money, stressful conditions, bureaucratic procedures, dictatorial bosses, or exploitation. However, next to sheltering people from sources of distress, the GB would also be challenging them to try out new things and to relentlessly advance beyond what they have now. In this way, the GB would provide the ideal environment for personal growth, self-actualization (Heylighen, 1992), and overall well-being.

Unlike the “rapture of the nerds”, with its flight from the material world into a Platonic Heaven, this “return to Eden” vision keeps us firmly with our feet down on earth. In the Earth’s restored “garden” state, we would enjoy its flowers, fruits and animals, and, more generally, our own sensations, bodies, activities, friends, lovers, and the natural world that brought them forth. Of course, the virtual world of the Global Brain would provide an astonishingly powerful extension to

our brains and senses. However, it will not replace the real world of matter, life, mind and society, which is immensely more complex, diverse, surprising and fascinating than any computational model imagined by Platonist nerds.

Once the GB (which of course includes us, humans) has tackled the more pressing challenges, we should expect it to expand its focus to domains that are still out of reach. The most obvious direction in which to look is outer space. However, the intuitive expectation that a global brain would expand into an interstellar, galactic, or even cosmic brain (Vidal, 2013) immediately runs into a hard limit: the finite speed of light implies that a signal sent to other stars would require decades, if not millennia, for an answer to return—a communication speed incompatible with the processing speed of a brain. A more likely expansion is into “inner space”: the microscopic realm of cells, organelles, molecules, atoms, elementary particles, and even the structure of space-time itself (Egan, 2002). Nanotechnology already promises to extend our powers of sensation and manipulation into the cellular domain, whose richness we have as yet hardly skimmed. But there is still “plenty of room at the bottom” for the GB to explore...

### **Perils on the road to the Singularity**

The previous section has probably suggested a too rosy picture of the transition towards a GB-dominated world. Technological advances, however beneficial in the long term, tend to create great stresses on society. First, new technologies need to adapt to the people that use them, a process that is everything but obvious. Second, people need to adapt to the uses they offer, something even more difficult. Moreover, the most beneficial uses typically result from *exaptation*, i.e. the discovery of functions or applications other than the ones originally intended (Andriani & Cohen, 2013; Dew, Sarasvathy, & Venkataraman, 2004). Vice versa, the originally intended applications often turn out to have more negative than positive effects.

This process of mutual adaptation and exaptation nearly always takes more time and effort than anticipated by the technology’s designers. One reason is that people are not rational—in the economic sense of making optimal decisions based on full information—but driven by habit, fashion and tradition; biases; shallow impressions; and deep emotions. This applies both to designers and to users. Another reason is that people and artifacts together form a highly complex and non-linear system. Such systems often exhibit erratic and counter-intuitive behaviors that make it impossible to predict the effects and side effects of a specific innovation. Let us then look in more detail at some of the most important things that can go wrong.

### **Cascading failures**

The non-linearity of the global system is exacerbated by what has been called *hyperconnectivity* (Fredette, Marom, Steiner, & Witters, 2012): nowadays nearly everyone and everything is connected to everything via global networks. This makes it easy for disturbances to propagate from node to node in the network, spreading across ever widening circles ever more quickly. Examples are

computer viruses or false information reaching across the globe in hours via the Internet, or real viruses spreading in days via intercontinental travel. Less obvious but equally dangerous are local breakdowns—such as a bank going broke or a power line that is overloaded—that precipitate similar breakdowns in the systems that directly depend on them. These in turn can precipitate further breakdowns, until a large-scale network collapses. Examples of such series of *cascading failures* are the 2008 financial collapse and various large-scale electricity blackouts (Dueñas-Osorio & Vemuru, 2009; Helbing, 2013).

The underlying mechanism is a positive feedback or vicious cycle that amplifies the disturbance more quickly than it can be contained. The infrastructure of the GB, which reduces friction and therefore facilitates propagation, increases such systemic risks (Heylighen, 2007a). On the other hand, a fully functioning GB should be able to sense problems, devise solutions, and intervene nearly immediately, thus nipping any potential crisis in the bud. Unfortunately, we are not yet at the stage that we can rely on such a superhuman distributed intelligence. In the meantime, we will have to develop more basic protection measures. These include:

- the equivalent of “firewalls” and artificially imposed friction to contain too fast spreading,
- sufficient redundancy so that the function of failing systems or components can be taken over by others,
- sufficient diversity in designs so that only a small proportion of systems are vulnerable to the a specific type of failure,
- large enough reserves or buffers to compensate for any large-scale shortages,
- an artificial “immune system” that would automatically detect potentially dangerous phenomena, and block them from spreading until they are proven safe,
- a better understanding and pre-emptive remediation of the most important vulnerabilities—e.g. by strengthening the nodes in the network that most other nodes (in)directly depend on (Farmer, 2013; Thurner & Poledna, 2013).

Without such precautions, our society runs the risk of suffering truly global crises and catastrophes (Helbing, 2013).

### Cyber war

Next to such risks for accidental damage, we need to consider the risks of intentional harm done to, or through, the network. This is the domain of activities that have been labeled *hacking*, *cyber crime*, *cyber terrorism*, and even *cyber war*. The principle is that some people or organizations exploit weaknesses in the systems connected via the Internet, in order to access secret data, steal money, wreak damage, or hinder services. One of the most spectacular examples is the Stuxnet worm, developed by US and Israeli intelligence agencies in order to sabotage Iranian nuclear plants by corrupting their software (Schmidt & Cohen, 2013).

These problems are not fundamentally different from the ones encountered before the advent of the net: there has always been an arms race between police and criminals, spies and counterspies, powers and rival powers. Whenever one of these parties appears to get the upper hand by developing

a new weapon or tactic, or by discovering a hole in the other's line of defense, the other party will eventually plug the hole or develop a counter-tactic.

The same will continue to happen on the Internet over the next decades. The main difference is that the Internet is more complex, more anonymous, and farther removed from the physical world in which bodily harm can occur. Therefore, these on-going skirmishes are likely to remain complicated and murky, often with no clear victims or perpetrators, while having relatively little effect on the general population (Schmidt & Cohen, 2013). Moreover, given the great economic interdependency of the major nations, there is little to gain, and much to lose by seriously damaging the other party's capabilities. More generally, the on-going spread of democracy and Enlightenment values, which is facilitated by the Internet, strongly reduces the willingness to engage in war, conflict or any kind of violence (Mueller, 2009; Pinker, 2011).

Therefore, most of these covert activities will remain limited to stealing confidential information, such as blueprints for new technologies or political strategies. This threat will gradually lose in significance as the "open access" ideology that drives the GB makes information more freely available (Heylighen, 2007b), and as better laws and technologies for the protection of privacy and personal identity are put into place. Thus, a "cyber war" is better than a "cold war", which is itself preferable to a traditional war. The only serious danger is that one of the pieces of malicious software would start spreading beyond its target, and thus harm the world as a whole. But this merely points us back to the issue of *cascading failures* and our list of proposed measures for containing the "viral" spread of disturbances.

### Psychological parasites

There is another old problem that will get a new life on the Internet, however this time with potentially more widespread harm: *psychological parasites*. Our brain is naturally attracted to certain exciting, calming or pleasure-inducing stimuli, such as sex, food, or praise. Therefore, it tends to seek more of such sensations, even in cases when they are not accompanied by any real benefit. The danger is that such stimuli can become *parasitic*: they get themselves multiplied by exploiting people's instinctive desires, without regard for their well-being. An example of such a mental parasite is a drug addiction: drug-induced sensations reproduce themselves by pushing the addict to seek the same stimulus again and again. There exists a wide variety of other addictions, including alcoholism, gambling, and sex addiction, in which people feel compelled to perpetually recreate the same powerful stimuli, even though this activity harms their health and social functioning.

When the stimuli are produced by electronic media, the result may be called *cyber addiction* (Beard, 2005; Chou & Ting, 2003; Young, 1998). This includes addictions to computer games, Internet use, social media, and smartphones. The danger of cyber addiction can only grow as software becomes more responsive, human-computer interfaces become more immersive, and ambient intelligence learns to better please its users (Heylighen et al., 2013). The result is that network users run a serious risk of becoming enslaved into a pattern of repetitive, sensation-seeking activity, only because this activity pushes all the right "buttons" in their brain at the right moments.

Three psychological mechanisms in particular appear powerful enough to drive the global spread of such psychological parasites: *flow*, *supernormal stimuli* and *mind viruses*.

*Flow* is the pleasurable state achieved when a person is so engulfed in an activity that s/he only wants to continue that activity, ignoring all other concerns (Nakamura & Csikszentmihalyi, 2002). The addictive quality of computer games is commonly explained by their ability to produce flow (Chen, 2007; Chou & Ting, 2003; Heylighen et al., 2013). Flow is achieved when challenges are presented that are matched to the person's skills in achieving them, and every action by the user receives an immediate feedback, so that the user is continuously stimulated to go further and further. Such flow-producing stimulation is easily implemented in what I have called a *mobilization system* (Heylighen et al., 2013). This is an ICT system that encourages and coordinates human action. Such technologies are intended to mobilize people for worthwhile objectives, such as educating themselves, exercising, or collaboratively solving problems. However, their potential power on the human mind is such that they can lead to both addiction and exploitation by political or commercial organizations. This would turn their users into unwitting slaves of the system (Heylighen et al., 2013).

Supernormal stimuli are stimuli more intense than those that occur in the natural world (Barrett, 2010). In most cases, the brain pays attention not so much to the absolute, but to the relative intensity of a stimulus: in how far it is stronger or weaker than other stimuli? In the competition between stimuli, the strongest one normally wins—even if it takes on absurd proportions. This can be illustrated by cartoon figures—such as Mickey Mouse—which tend to have impossibly large heads and eyes supported by ridiculously short legs. That is because faces and eyes are intrinsically more interesting to the brain. Therefore, artists have learned to exaggerate those features in order to attract the attention. Similar tricks are used in a variety of domains, including advertising, computer games, junk food (which contains unnaturally high concentrations of sugar, fat, salt, and calories), fashion models (who tend to have abnormally long and thin bodies), movies (which show ever more extreme special effects and violence), etc.

In a society overloaded with information, interruption and stimulation, attention is scarce (Heylighen et al., 2013). This creates competition for the attention of the public. The result is an arms race towards ever more powerful stimuli, which deviate ever further from what exists in nature. This leads to unhealthy addictions (e.g. to junk food, violent movies, or Internet porn) and to stress-inducing overstimulation. The even greater danger is that people would lose touch with reality, so that they develop completely unrealistic expectations about what a sex partner is supposed to look like, how much violent crime there is in the streets, or what constitutes a healthy diet. This makes them intrinsically dissatisfied, because the things they can get in the real world seem pale in comparison with the things they see in the virtual world of games and advertising. Moreover, it makes them lose their sense of balance and judgment. Thus, they are more likely to lapse into extreme behaviors, such as joining a cult, becoming a suicide terrorist, or going on a shooting spree to kill dozens of bystanders—because these do not seem so extreme anymore compared to the standards of their video games or horror movies.

*Mind viruses* (Brodie, 1996) can be defined as parasitic memes. A *meme* is an idea, belief or behavior that spreads across society by being transmitted from person to person (Heylighen &

Chielens, 2008). Successful memes tend to exhibit characteristics such as plausibility, simplicity, novelty, usefulness, emotional impact, and ease of communication. As long as a meme provides valuable information, its propagation benefits society. However, parasitic memes mimic the characteristics of beneficial memes in order to spread more easily, while providing information that is worthless, wrong, or even dangerously misleading (Heylighen & Chielens, 2008). Examples include chain letters, false rumors, urban legends, hate speech, conspiracy theories, superstitions, extremist ideologies, and various fundamentalist and irrational beliefs.

Such memes are particularly dangerous to individuals who already have lost their sense of reality by their immersion in supernormal and flow-producing stimulation environments, and who thus are ready to embrace the false promises of a mind virus. Because they spread across communities while indoctrinating their carriers, mind viruses have an even greater potential to create damage. For example, they may recruit a worldwide group of people into an absurd and destructive enterprise, such as an outbreak of gratuitous rioting, a mass suicide, a terrorist plot, or even a genocide. The ubiquitous network enhances their powers of spreading and mobilization, thus increasing the danger.

In principle, an omnibenevolent GB would want to prevent such disastrous outcomes. However, omnibenevolence cannot be assumed to be in-built or preprogrammed by some rational intelligence. It is rather the result of self-organization in which technologies, ideas and institutions that seem to better satisfy the needs of their users are preferentially retained, propagated and further developed (Heylighen, 2013). But since humans have bounded rationality, they cannot foresee the long-term effects of such innovations, and may adopt a solution that satisfies their short-term need for stimulation or group belonging, while failing to see that it is detrimental to their long-term needs for health and constructive engagement with society. Therefore, it may take quite a while before society has become fully aware of the dangers of such parasitic stimuli, allowing them to cause a lot of harm in the mean time.

To avoid such damage, we will need to develop precautionary measures similar to the ones Helbing (2013) advocates for systemic risks, but this time including the social and emotional characteristics of people and communities in addition to the dynamics of cascading disturbances. This should allow us to build another “global immune system” that would be able to anticipate, detect, quarantine, and if necessary neutralize psychological parasites. For this, we could rely on the criteria that characterize successful memes (Heylighen & Chielens, 2008) and mobilization systems (Heylighen et al., 2013), coupled with a realistic simulation of how messages spread across the global brain network (Heylighen, Busseniers, Veitas, Vidal, & Weinbaum, 2012). Armed with such knowledge, we may be able to recognize potentially epidemic parasites before they have spread, forecast their path of propagation, and thus prevent their spreading by “inoculating” the people likely to be infected. This can be done by sending them persuasive warnings or refutations about the propagating rumor, terrorist propaganda, or addictive system before they have been exposed to them.

In addition, we may need to develop laws and institutions that prevent the willful creation and propagation of psychological parasites. Already existing examples are laws that prohibit hate speech, drug trafficking, or membership of terrorist groups, but these are too specific to regulate more general threats. Such a legal framework would make it more difficult to exploit unsuspecting people

by bombarding them with dangerous, addictive, and self-reinforcing stimuli (as the tobacco industry has been doing for decades, and the junk food industry is still doing).

### Future Shock

As Toffler (1970) already discussed in great detail, having to adapt to new circumstances is intrinsically stressful. Therefore, people are inclined to resist too fast or too radical changes. If they cannot, they will be left confused, frightened and disorientated—a condition that Toffler called “future shock”. The reason that I suggested that sociotechnological acceleration may already have reached its peak is that culture, politics and society seem to evolve more slowly now than they did in the 1960s, the decade that inspired Toffler’s analysis. For example, musical styles, clothes fashions, literary genres, intellectual theories, and political ideologies seem to be hardly different from what they were 20 years ago, in contrast to the revolutionary changes seen in the 1960s and 1970s.

A plausible explanation for this slowdown is that it compensates for the acceleration in ICT and its impact on people’s lives over the last two decades. When you need to learn radically new technologies, applications and ways of communicating every few years, you will have less energy to explore new styles of music or literature, or to rebel against your parents’ worldview. While there seems to be no slowdown in the creation of new technologies, their speed of adoption by the wider public may well be close to its limit. The inertia is even greater in political, economical and legal institutions, which tend to lag far behind in adapting to the social implications of ICT. Thus, they are ill prepared to deal with the perils and promises we surveyed. This creates great stresses on both society and its members, making them even more vulnerable to psychological parasites and cascading failures.

The most likely negative outcome is a conservative backlash, in which people or political regimes desire to turn back the clock to an idealized past and its simple truths, thus trying to repel the waves of change brought by globalization and technological innovation. The problem, of course, is that these past solutions are even less adapted to the present turmoil. They can thus only exacerbate the problem. In its extreme form, such a backlash can take the form of a fundamentalist uprising, ending in either on-going violent conflict and terrorism, or a rigid, totalitarian regime. On a larger scale, this could give rise to wars between religious or ideological blocks, and a drastic reduction in freedom and creativity. This would obviously constitute a serious setback on the road to the GB, creating a lot of needless suffering.

In a less extreme form, a conservative backlash will merely hinder further advances. This is not necessarily a bad outcome, as the resulting slowdown may help people to psychologically recover from their “future shock”, while giving institutions the time to effectively assimilate some of the changes. However, the preferable strategy is to go for a controlled “soft landing”, in which the less important but relatively stressful changes (e.g. new interfaces, operating systems, or bureaucratic procedures) are artificially restrained, in order to keep the most constructive ones (e.g. better systems for collective intelligence, global governance, online education...) up to speed.

To further cushion the shock, Toffler (1970) proposed to maintain “enclaves of the past”, in which people who are not ready to embrace the brave new world of the emerging GB may continue

to live more or less as they used to, limiting themselves consciously to a simpler life. Two examples at the opposite poles illustrate the variety of possible strategies:

1. the Pennsylvania Amish have not essentially changed the technologies they use since the 18<sup>th</sup> century;
2. ICT systems that change their interface often have the option for the user to maintain the old interface for a couple of months.

Ideally, there should be a whole range of options in between so that people can adapt to innovations at their own pace. In particular, new technologies should as much as possible be backward compatible, so that people using older systems would not a priori be excluded. This should reduce both overall stress and the likeliness of a radical backlash.

### Loss of human capabilities

We noted that it is in the GB's interest to maximize human capabilities. However, in the short term there is a serious risk that certain capabilities will get lost, at least in some populations. Technology is intended to make life easier, thus reducing physical and intellectual demands on its users. GB technologies in particular should be able to tackle nearly any challenge with a minimum effort required from individuals. For example, someone who needs to solve a problem could simply request the GB's assistance and blindly follow its recommendations—while being carried to the right place by the robotic vehicle the GB has dispatched.

This makes it very tempting to adopt an attitude of passivity, in which individuals avoid any challenge more difficult than the ones they feel perfectly comfortable with. Without challenges to stimulate the development of mental and physical skills, the supporting organs and circuits tend to atrophy—like the muscles that waste away when a person has been staying too long in a hospital bed. This is the well-known principle of “*use it or lose it*”. It applies in particular to older or sick people, who may have more difficulty than before to perform a certain task, and therefore stop doing it altogether—even after their health has improved again. The resulting degeneration of abilities seems irreversible, as further loss of function triggers further avoidance, and so forth. More generally, commentators lament the general loss of physical fitness in a population that seems ever less inclined to walk, run, climb, carry or perform any natural exertion, while remaining attached to a screen the whole day.

Loss of physical ability is not a recent phenomenon, though. Analysis of the remains of our Paleolithic ancestors indicates that they were taller, stronger, and healthier than their descendants (Diamond, 1987; Mummert, Esche, Robinson, & Armelagos, 2011). Even their brains were about 10% larger than ours (McAuliffe, 2010), suggesting a higher intelligence. A plausible explanation is that these hunter-gatherers had a physically and intellectually more challenging life than the subsequent farmers and industrial workers, since the latter could rely on society and technology (including agriculture) to satisfy the bulk of their needs. Another likely reason is that their lifestyle and diet were more in tune with what their organism was adapted to (De Vany, 2010).

Both differences are exacerbated in our present information age:

1. we rely even more on technology and society to solve all our problems;
2. we tend to eat and live in a way that is even further removed from what nature intended.

This unnatural lifestyle leads to an all too recognizable list of “diseases of civilization”: obesity, diabetes, metabolic syndrome, cardiovascular diseases, cancer, depression, allergies, autoimmune diseases, ADHD, etc. (Carrera-Bastos, Fontes-Villalba, O’Keefe, Lindeberg, & Cordain, 2011). The consequences seem particularly dangerous for our children (Palmer, 2010), who have never even experienced the more traditional lifestyle before the advent of electronic media, and who tend to develop problems like diabetes and obesity already from an early age. Extrapolating from this epidemic, some analysts suggest that for the first time in a millennium the life expectancy of the next generations may be lower than the one of the present generation (Olshansky et al., 2005).

The attractiveness of such unnatural lifestyles is at least in part to blame on the supernormal stimuli, such as junk food, game addiction, television violence etc., that they offer. Lack of real challenges in our daily life is compensated by the overload of artificial stimuli we experience while watching a horror movie, playing a virtual reality adventure game, or visiting a porn site. While it is clear that this can only weaken our physical abilities, the question remains open as to its effect on our mental abilities. Observers tend to be split, some pointing towards an increasingly shallow comprehension and loss of the ability to concentrate (Carr, 2011), others towards the cognitive stimulation provided by television and computer games (Johnson, 2006). The overall effect is likely to depend on:

- (1) what kind of electronic media are used;
- (2) which functions of the brain are being stimulated.

Concerning point (1), more complex, interactive and challenging environments will obviously have a more beneficial effect than more shallow, passive, and supernormal stimulation. The choice of stimulation will depend on the individual, the social environment, and the available media. Typically, people from a lower social and educational background will gravitate towards the less intellectually challenging media, while the opposite will happen for those from a higher background. This will exaggerate the existing intellectual differences, boosting the capabilities of the smart people, while “dumbing down” the not so smart. Moreover, people from poorer backgrounds tend to lead a more unhealthy and stressful life, making them particularly inclined to avoid any additional challenges. The danger is that this will further amplify social inequalities, creating an “underclass” of people who simply cannot grasp the complexities of a high-tech society.

As to point (2), it seems likely that even the most complex and intellectually challenging virtual media will tend to boost certain brain functions, while neglecting others. Computer use typically demands symbolic skills, such as language, math, and logical reasoning, and specialized visual skills. On the other hand, tactile, motor, social and emotional forms of intelligence are likely to develop better by interacting with the real world. As people—and children in particular—tend to spend more time in the virtual world, they are likely to gain in symbolic intelligence, but lose in “natural” or intuitive intelligence. While a boost in symbolic skills may be necessary to fully take advantage of ICT applications, in the longer term the effect may be counterproductive. Indeed, computers are intrinsically better at symbolic reasoning than humans, but worse at embodied and intuitive reasoning. For the GB, the ideal balance is to delegate tasks of the symbolic type to

computers, and the other tasks to humans. But that would require that humans maximally develop their intuitive and embodied skills.

In the longer term, the GB is likely to get the balance right by creating an environment that challenges people both physically and mentally to develop themselves as much as they can. Moreover, the ongoing scientific and technological advances, in domains such as medicine, sport, drugs, intelligence amplification, genetic manipulation, nanotechnology and cybernetic implants, is likely to boost the capabilities of individuals to a “transhuman” level—albeit that just getting back to the level we had in the Paleolithic would already be a remarkable achievement...

In the meantime, we will need to make sure that physical and intellectual passivity, supernormal stimuli and alienation from our natural environment do not get out of control, especially among the least educated. Targeted education and stimulation, supported by mobilization systems (Heylighen et al., 2013), should be able to achieve these objectives. However, until they take full effect we are likely to witness an enormous waste of human life, health, fitness, well-being and intellectual capabilities, just because technology lures people into choosing the path of least resistance (and strongest stimulation)...

## Conclusion

I have sketched the outlines of a scenario in which technological acceleration culminates in the self-organization of a distributed intelligence out of the global ICT network. This intelligent network would play the role of a brain for the global superorganism—the living system consisting of all people, their artifacts, and the ecosystems they depend on. As such, this brain would be responsible for tackling all challenges that may confront both the planet as a whole and its human components. By harnessing the collective intelligence of all people, while complementing it with the data processing capabilities of an immense array of computers, sensors, robots, and other technologies, this Global Brain (GB) would achieve a level of superhuman capabilities that are as yet difficult to imagine.

To better convey their significance, I have invoked a pragmatic version of the classic list of “divine” attributes:

- *omniscience*: being able to provide the knowledge needed to tackle any real challenge;
- *omnipresence*: being able to sense, inform and act anytime anywhere on Earth;
- *omnipotence*: being able to produce any good, service or coordinated action at a minimal cost;
- *omnibenevolence*: using these abilities to produce the greatest happiness for the greatest number of people.

These attributes are not so much the characteristics of an idealized, God-like entity, but the features representing the main selective pressures on socio-technological evolution: both designers and users of technology would like their systems to become more intelligent, more knowledgeable, more widely available, more powerful, and more beneficent in their effects. This means that in the longer

term (after adjusting for misjudgments, side effects and exaptations) systems and components that satisfy these criteria better should replace those that do so less well.

Moreover, I have surveyed a number of existing trends and technologies that illustrate how most of these capabilities may be reached within the next few decades. For example, Wikipedia already provides the bulk of human knowledge to anyone interested, location-aware smartphones and wireless sensors are spreading rapidly across the globe, 3D printers are in principle capable of realizing any design downloaded from the Internet (Lipson & Kurman, 2013), while violence and conflict around the world have reached historically low levels (Pinker, 2011). Thus, Kurzweil's (2005) prediction of a technological singularity around the year 2045 may not be so far off the mark—at least if we interpret this singularity as the full deployment of a Global Brain (Heylighen, 2008), and not as the creation of a superhuman, autonomous AGI program. After that transition, we may experience a true *return to Eden*—an idyllic state of abundance, peace and well-being, in which all serious threats have been tackled and people can fully dedicate themselves to further creative endeavors.

Before we reach such a utopian ideal, however, the road to the GB threatens to be long and arduous. While technological innovation tends to be beneficial in the long run, in the short run it is accompanied by side effects and dangers that are difficult to predict and to control. One fundamental problem brought about by the network revolution is the general increase in connectivity and reduction in friction. This makes it easier for various phenomena, positive as well as negative, to propagate. Thus, they may spread so quickly across the globe that they run out of control. Particularly dangerous are cascading failures, such as collapsing stocks triggering the collapse of other stocks, and viral phenomena, such as rumors, fads, religious cults, computer viruses, and infectious diseases, which propagate nearly unhindered across the global communication network. To avoid major crises caused by such self-amplifying phenomena, we will need to set in place a range of precautionary measures. These include highly diverse and redundant reserve capacities, “firewalls”, and a “global immune system” that would quickly recognize and immobilize potentially dangerous “virals”.

Not just the communication network, but our human psychology is particularly vulnerable to parasitic phenomena and negative side effects. As technology lures us further away from the natural environment to which our bodies and minds are adapted, we become more inclined to accept some of its seemingly attractive, but ultimately destructive, offerings. These include drugs, junk food, addictive games and virtual environments, sensational videos, and the allure of spending the whole day in the sofa jacked into the network, without moving or experiencing fresh air and sunlight. We may even be tempted to stop tackling any real challenges ourselves, and start leading the life of a decadent aristocrat—with the GB in the role of the omnipresent butler that keeps things going and cleans up all the mess.

All this, coupled with the intrinsic stress of having to adapt to accelerating change, is likely to severely strain our mental and physical health, fitness and happiness. Given that certain groups are particularly vulnerable to the resulting loss of capabilities, this may well result in a class of deeply frustrated, maladapted and alienated people, who are easy prey for extremist, reactionary and fundamentalist movements and ideas. Recommended precautions here are measures to control stress

and the speed of change, the creation of a different type of “global immune system” that would preempt the spread of psychological parasites, and the deployment of mobilization systems that challenge people to develop their physical and mental abilities in the most healthy manner.

In sum, we definitively live in the most interesting of times. While the potential for crises and catastrophes remains great, none of these seems serious enough to wipe out modern civilization, given the amazing capabilities of adaptation it has shown up to now. Therefore, I remain fundamentally optimistic. Over the past centuries, technological innovation has spectacularly improved the human condition (Heylighen & Bernheim, 2000), and the latest global statistics on human development show no sign of any slowdown or reversal in this long-term trend (Klugman, 2010). Moreover, given my reasons to discard any self-amplifying AI explosion and to believe in the universal benevolence of the GB, I have no fear for the many scenarios in which a superintelligent computer system would enslave or destroy humanity (Chalmers, 2010; Eden et al., 2013). Such scenarios have even prompted intelligent thinkers to call for a moratorium on AI research—a measure that I consider absolutely counterproductive—or to “lock up” any AI inside a virtual world where it could not affect the real world (Chalmers, 2010)—a measure that would guarantee that the AI would lack any practical intelligence.

I think it is time to get back with our two feet firmly on the ground, and continue investigating and improving the real world with all its complexities, interactions and nuances, instead of getting carried away by abstract speculations grounded in bits and bytes rather than in flesh and earth. In that way, we may truly be able to build a new Eden here on our magnificent planet, and thus undo most of the negative side effects that have tended to accompany civilization and technology up to now.

## Author Bio

I am a research professor at the Free University of Brussels (Vrije Universiteit Brussel), where I direct the transdisciplinary research group on "Evolution, Complexity and Cognition" and the Global Brain Institute. I am affiliated with the Center "Leo Apostel", and an editor of the Principia Cybernetica Project. The main focus of my research is the evolution of complexity: how do higher forms of organization originate and develop? How do systems self-organize, adapt and achieve some form of cognition? I have worked in particular on the development of collective intelligence or distributed cognition, and its application to the emerging "global brain".

## References

- Akyildiz, I. F., Wang, X., & Wang, W. (2005). Wireless mesh networks: a survey. *Computer Networks*, 47(4), 445–487. doi:10.1016/j.comnet.2004.12.001
- Anderson, C. (2012). *Makers: the new industrial revolution*. Random House.
- Andriani, P., & Cohen, J. (2013). From exaptation to radical niche construction in biological and technological complex systems. *Complexity*, 18(5), 7–14. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1002/cplx.21450/full>

- Ashby, W. R. (1958). Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1(2), 83–99.
- Ashby, W. R. (1964). *An introduction to cybernetics*. Methuen London.
- Ashby, W. R. (1972). Setting goals in cybernetic systems. In H. W. Robinson & D. E. Knight (Eds.), *Cybernetics, Artificial Intelligence and Ecology* (pp. 33–44). New York: Spartan Books.
- Ashby, W. R. (1981). *Mechanisms of Intelligence*. (R. Conant, Ed.). Intersystems Publications. Retrieved from <http://books.google.be/books?hl=en&lr=&id=vbcEczxQgpkC&oi=fnd&pg=PR1&dq=ashby+intelligence+amplifier&ots=6K4fWa6rOA&sig=5mOvDAnVhDDjbXHjmbY1jVyx5Y>
- Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. *Computer Networks*, 54(15), 2787–2805. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1389128610001568>
- Barfield, W., & Caudell, T. (2001). *Fundamentals of wearable computers and augmented reality*. Routledge. Retrieved from <http://books.google.be/books?hl=en&lr=&id=-1w4367mu3QC&oi=fnd&pg=PR11&dq=wearable+computers&ots=xqPLZGY1Bq&sig=txozDYG538Tp15mP4r8wy1uvNY0>
- Barrett, D. (2010). *Supernormal stimuli: how primal urges overran their evolutionary purpose*. W.W. Norton & Co.
- Barrow, J. D. (1998). *Impossibility: The limits of science and the science of limits*. Oxford University Press, USA.
- Beard, K. W. (2005). Internet addiction: a review of current assessment techniques and potential assessment questions. *CyberPsychology & Behavior*, 8(1), 7–14. Retrieved from <http://online.liebertpub.com/doi/pdf/10.1089/cpb.2005.8.7>
- Berners-Lee, T., & Fischetti, M. (1999). *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Harper San Francisco. Retrieved from <http://portal.acm.org/citation.cfm?id=554813>
- Bernheim, J. (1999). The Cognitive Revolution and 21st Century Enlightenment: towards a progressive world view. *Science, Technology and Social Change*, Einstein meets Magritte (p. 63). Kluwer.
- Bernstein, A., Klein, M., & Malone, T. W. (2012). Programming the Global Brain. *Communications of the ACM*, 55(5), 1. Retrieved from <http://cci.mit.edu/publications/CCIwp2011-04.pdf>
- Bohn, J., Coroamă, V., Langheinrich, M., Mattern, F., & Rohs, M. (2005). Social, economic, and ethical implications of ambient intelligence and ubiquitous computing. *Ambient intelligence* (pp. 5–29). Springer. Retrieved from [http://link.springer.com/chapter/10.1007/3-540-27139-2\\_2](http://link.springer.com/chapter/10.1007/3-540-27139-2_2)
- Brodie, R. (1996). *Virus of the Mind: The New Science of the Meme*. Integral Pr.
- Carr, N. G. (2011). *The shallows: what the Internet is doing to our brains*. New York: W.W. Norton.
- Carrera-Bastos, P., Fontes-Villalba, M., O’Keefe, J. H., Lindeberg, S., & Cordain, L. (2011). The western diet and lifestyle and diseases of civilization. *Research Reports in Clinical Cardiology*, 15. doi:10.2147/RRCC.S16919
- Chalmers, D. (2010). The Singularity: A philosophical analysis. *Journal of Consciousness Studies*, 17, 9(10), 7–65.
- Charlton, B. (2002). What is the Meaning of Life? Animism, Generalised Anthropomorphism and Social Intelligence. Retrieved from <http://www.hedweb.com/bgcharlton/meaning-of-life.html>
- Chen, J. (2007). Flow in games (and everything else). *Communications of the ACM*, 50(4), 31–34.
- Chou, T.-J., & Ting, C.-C. (2003). The role of flow experience in cyber-game addiction. *CyberPsychology & Behavior*, 6(6), 663–675. Retrieved from <http://online.liebertpub.com/doi/abs/10.1089/109493103322725469?2>
- Clark, A. (1998). Embodied, situated, and distributed cognition. *A companion to cognitive science*, 506–517.

- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Corning, P. (2003). *Nature's magic: Synergy in evolution and the fate of humankind*. Cambridge University Press Cambridge, UK.
- Cullen, J., & Bryman, A. (1988). The knowledge acquisition bottleneck: time for reassessment? *Expert Systems*, 5(3), 216–225. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0394.1988.tb00065.x/abstract>
- Dawkins, R. (1999). *The extended phenotype: The long reach of the gene* (Revised.). Oxford University Press, USA.
- De Alfaro, L., Kulshreshtha, A., Pye, I., & Adler, B. T. (2011). Reputation systems for open collaboration. *Communications of the ACM*, 54(8), 81–87. Retrieved from <http://dl.acm.org/citation.cfm?id=1978560>
- De Rosnay, J. (2000). *The Symbiotic Man: A new understanding of the organization of life and a vision of the future*. McGraw-Hill. Retrieved from <http://pespmc1.vub.ac.be/books/DeRosnay.TheSymbioticMan.pdf>
- De Vany, A. (2010). *The New Evolution Diet: What Our Paleolithic Ancestors Can Teach Us about Weight Loss, Fitness, and Aging*. Rodale Books.
- Dew, N., Sarasvathy, S. D., & Venkataraman, S. (2004). The economic implications of exaptation. *Journal of Evolutionary Economics*, 14(1), 69–84. Retrieved from <http://link.springer.com/article/10.1007/s00191-003-0180-x>
- Diamandis, P. H., & Kotler, S. (2012). *Abundance: The Future Is Better Than You Think*. Free Press.
- Diamond, J. (1987). The Worst Mistake In The History Of The Human Race. *Discover*, May, 64–66.
- Dressler, F. (2008). A study of self-organization mechanisms in ad hoc and sensor networks. *Computer Communications*, 31(13), 3018–3029.
- Drexler, E. K. (2013). *Radical Abundance: How a Revolution in Nanotechnology Will Change Civilization*. PublicAffairs.
- Dueñas-Osorio, L., & Vemuru, S. M. (2009). Cascading failures in complex infrastructure systems. *Structural Safety*, 31(2), 157–167. doi:10.1016/j.strusafe.2008.06.007
- Eden, A. H., Raker, J. H. S., Moor, J. H., & Steinhart, E. H. (2013). *Singularity Hypotheses: A Scientific and Philosophical Assessment*. Springer-Verlag New York Incorporated. Retrieved from <http://www.springer.com/engineering/computational+intelligence+and+complexity/book/978-3-642-32559-5>
- Egan, G. (2002). *Schild's Ladder* (1st ed.). Eos.
- Farmer, J. D. (2013). *The challenge of modeling systemic risk*. University of Oxford. Retrieved from <http://www.risk.jbs.cam.ac.uk/news/events/annualmeetings/downloads/pdfs/2013/farmerj.pdf>
- Fredette, J., Marom, R., Steiner, K., & Witters, L. (2012). The promise and peril of hyperconnectivity for organizations and societies. *The Global Information Technology Report: living in a hyperconnected world* (pp. 113–119). World Economic Forum. Retrieved from [http://www3.weforum.org/docs/GITR/2012/GITR\\_Chapter1.10\\_2012.pdf](http://www3.weforum.org/docs/GITR/2012/GITR_Chapter1.10_2012.pdf)
- Gao, H., Barbier, G., & Goolsby, R. (2011). Harnessing the crowdsourcing power of social media for disaster relief. *Intelligent Systems, IEEE*, 26(3), 10–14.
- Gershenson, C., & Heylighen, F. (2005). How can we think complex? In K. Richardson (Ed.), *Managing organizational complexity: philosophy, theory and application* (pp. 47–61). Information Age Publishing. Retrieved from <http://pcp.vub.ac.be/Papers/ThinkingComplex.pdf>
- Goertzel, B. (2002). *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers.

- Google Glass. (2013, July 16). *Wikipedia, the free encyclopedia*. Retrieved from [http://en.wikipedia.org/w/index.php?title=Google\\_Glass&oldid=564433859](http://en.wikipedia.org/w/index.php?title=Google_Glass&oldid=564433859)
- Gray, P. (2009). Play as a foundation for hunter-gatherer social existence. *American Journal of Play*, 4, 476–522.
- He, B., Gao, S., Yuan, H., & Wolpaw, J. R. (2013). Brain–Computer Interfaces. In B. He (Ed.), *Neural Engineering* (pp. 87–151). Boston, MA: Springer US. Retrieved from [http://www.springerlink.com/index/10.1007/978-1-4614-5227-0\\_2](http://www.springerlink.com/index/10.1007/978-1-4614-5227-0_2)
- Helbing, D. (2013). Globally networked risks and how to respond. *Nature*, 497(7447), 51–59. Retrieved from <http://www.nature.com/nature/journal/v497/n7447/abs/nature12047.html>
- Helbing, D., Bishop, S., Conte, R., Lukowicz, P., & McCarthy, J. B. (2012). FuturICT: Participatory computing to understand and manage our complex world in a more sustainable and resilient way. *European Physical Journal Special Topics*, 214, 11–39. Retrieved from <http://adsabs.harvard.edu/abs/2012EPJST.214...11H>
- Heylighen, F. (1992). A cognitive-systemic reconstruction of Maslow’s theory of self-actualization. *Behavioral Science*, 37(1), 39–58. doi:10.1002/bs.3830370105
- Heylighen, F. (1999). Advantages and limitations of formal expression. *Foundations of Science*, 4(1), 25–56. doi:10.1023/A:1009686703349
- Heylighen, F. (2000). Evolutionary Transitions: how do levels of complexity emerge? *COMPLEXITY*, 6(1), 53–57.
- Heylighen, F. (2002). The global brain as a new utopia. *Zukunftsfiguren*. Suhrkamp, Frankfurt. Retrieved from <http://pespmc1.vub.ac.be/papers/GB-Utopia.pdf>
- Heylighen, F. (2006). Mediator Evolution: a general scenario for the origin of dynamical hierarchies. *Worldviews, Science and Us*. (Singapore: World Scientific).
- Heylighen, F. (2007a). The Global Superorganism: an evolutionary-cybernetic model of the emerging network society. *Social Evolution & History*, 6(1), 58–119. Retrieved from <http://pcp.vub.ac.be/papers/Superorganism.pdf>
- Heylighen, F. (2007b). Why is Open Access Development so Successful? Stigmergic organization and the economics of information. In B. Lutterbeck, M. Baerwolff & R. A. Gehring (Ed.), *Open Source Jahrbuch 2007* (pp. 165–180). Lehmanns Media. Retrieved from <http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf>
- Heylighen, F. (2008). Accelerating socio-technological evolution: from ephemeralization and stigmergy to the global brain. *Globalization as evolutionary process: modeling global change*, Rethinking globalizations (p. 284). Routledge. Retrieved from <http://pcp.vub.ac.be/papers/AcceleratingEvolution.pdf>
- Heylighen, F. (2011). Conceptions of a Global Brain: an historical review. *Evolution: Cosmic, Biological, and Social*, eds. Grinin, L. E., Carneiro, R. L., Korotayev A. V., Spier F. (pp. 274 – 289). Uchitel Publishing. Retrieved from <http://pcp.vub.ac.be/papers/GB-conceptions-Rodrigue.pdf>
- Heylighen, F. (2012a). A brain in a vat cannot break out: why the singularity must be extended, embedded and embodied. *Journal of Consciousness Studies*, 19(1-2), 126–142. Retrieved from <http://pcp.vub.ac.be/Papers/Singularity-Reply2Chalmers.pdf>
- Heylighen, F. (2012b). *Challenge Propagation: a new paradigm for modeling distributed intelligence* ( No. 2012-01). GBI Working Papers. Brussels, Belgium. Retrieved from <http://pcp.vub.ac.be/papers/ChallengePropagation.pdf>
- Heylighen, F. (2012c). *A Tale of Challenge, Adventure and Mystery: towards an agent-based unification of narrative and scientific models of behavior* (Working papers No. 2012-06). ECCO. Brussels, Belgium. Retrieved from <http://pcp.vub.ac.be/papers/TaleofAdventure.pdf>

- Heylighen, F. (2013). *Distributed Intelligence Technologies: a survey of present and future applications of the Global Brain* ( No. 2013-02). GBI Working Paper. Retrieved from <http://pespmc1.vub.ac.be/Papers/GB-applications-survey.pdf>
- Heylighen, F., & Bernheim, J. (2000). Global Progress I: Empirical Evidence for ongoing Increase in Quality-of-life. *Journal of Happiness Studies*, 1(3), 323–349. doi:10.1023/A:1010099928894
- Heylighen, F., & Bollen, J. (1996). The World-Wide Web as a Super-Brain: from metaphor to model. *Cybernetics and Systems '96. R. Trappl (Ed.). Austrian Society For Cybernetics.*
- Heylighen, F., & Bollen, J. (2002). Hebbian Algorithms for a Digital Library Recommendation System. *International Conference on Parallel Processing* (p. 439). Los Alamitos, CA, USA: IEEE Computer Society. doi:10.1109/ICPPW.2002.1039763
- Heylighen, F., Busseniers, E., Veitas, V., Vidal, C., & Weinbaum, D. R. (2012). *Foundations for a Mathematical Model of the Global Brain: architecture, components, and specifications* (GBI Working Papers No. 2012-05). GBI Working Papers. Retrieved from <http://pespmc1.vub.ac.be/papers/TowardsGB-model.pdf>
- Heylighen, F., & Chielens, K. (2008). Evolution of culture, memetics. (R. A. Meyers, Ed.) *Encyclopedia of Complexity and Systems Science*. Springer. Retrieved from <http://pespmc1.vub.ac.be/Papers/Memetics-Springer.pdf>
- Heylighen, F., & Joslyn, C. (2003). Cybernetics and Second-Order Cybernetics. In Robert A. Meyers (Ed.), *Encyclopedia of Physical Science and Technology (Third Edition)* (3rd ed., Vol. 4, pp. 155–169). New York: Academic Press. Retrieved from <http://pespmc1.vub.ac.be/Papers/cybernetics-EPST.pdf>
- Heylighen, F., Kostov, I., & Kiemen, M. (2013). Mobilization Systems: technologies for motivating and coordinating human action. In Peters M. A., Besley T. and Araya D. (Ed.), *The New Development Paradigm: Education, Knowledge Economy and Digital Futures*. Routledge. Retrieved from <http://pcp.vub.ac.be/Papers/MobilizationSystems.pdf>
- Hoffman, J., & Rosenkrantz, G. (2012). Omnipotence. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2012.). Retrieved from <http://plato.stanford.edu/archives/spr2012/entries/omnipotence/>
- Johnson, S. (2006). *Everything Bad is Good for You: How Popular Culture is Making Us Smarter*. Penguin Books Limited.
- Klugman, J. (2010). *Human development report 2010: The real wealth of nations*. New York: UNDP.
- Kurzweil, R. (2005). *The singularity is near*. Penguin books.
- LessWrong. (2013). Complexity of value. *LessWrong*. Wiki. Retrieved from [http://wiki.lesswrong.com/wiki/Complexity\\_of\\_value](http://wiki.lesswrong.com/wiki/Complexity_of_value)
- Lipson, H., & Kurman, M. (2013). *Fabricated: The New World of 3D Printing*. John Wiley & Sons.
- Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G., & Hartmann, B. (2011). Design lessons from the fastest q&a site in the west. *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11 (pp. 2857–2866). New York, NY, USA: ACM. doi:10.1145/1978942.1979366
- Mann, W. E. (1975). The Divine Attributes. *American Philosophical Quarterly*, 12(2), 151–159. Retrieved from <http://www.jstor.org/stable/10.2307/20009569>
- Mayer-Kress, G., & Barczys, C. (1995). The global brain as an emergent structure from the Worldwide Computing Network, and its implications for modeling. *The information society*, 11(1), 1–27.
- Maynard Smith, J., & Szathmáry, E. (1997). *The major transitions in evolution*. Oxford University Press, USA.
- McAuliffe, K. (2010). If Modern Humans Are So Smart, Why Are Our Brains Shrinking? *Discover Magazine*, (September). Retrieved from <http://discovermagazine.com/2010/sep/25-modern-humans-smart-why-brain-shrinking>

- Michael, D. R., & Chen, S. L. (2005). *Serious games: Games that educate, train, and inform*. Muska & Lipman/Premier-Trade.
- Miller, J. G. (1995). *Living systems*. University Press of Colorado.
- Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems: an introduction to computational models of social life*. Princeton University Press.
- Modis, T. (2012). Why the Singularity Cannot Happen. In A. H. Eden, J. H. Moor, J. H. Søraker, & E. H. Steinhart (Eds.), *Singularity Hypotheses* (pp. 311–346). Springer. Retrieved from [http://link.springer.com/chapter/10.1007/978-3-642-32560-1\\_16](http://link.springer.com/chapter/10.1007/978-3-642-32560-1_16)
- Muehlhauser, L., & Helm, L. (2012). The Singularity and Machine Ethics. In A. H. Eden, J. H. Moor, J. H. Søraker, & E. H. Steinhart (Eds.), *Singularity Hypotheses*, The Frontiers Collection (pp. 101–126). Springer Berlin Heidelberg. Retrieved from [http://link.springer.com/chapter/10.1007/978-3-642-32560-1\\_6](http://link.springer.com/chapter/10.1007/978-3-642-32560-1_6)
- Mueller, J. (2009). War has almost ceased to exist: An assessment. *Political Science Quarterly*, 124(2), 297–321. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1002/j.1538-165X.2009.tb00650.x/abstract>
- Mummert, A., Esche, E., Robinson, J., & Armelagos, G. J. (2011). Stature and robusticity during the agricultural transition: Evidence from the bioarchaeological record. *Economics & Human Biology*, 9(3), 284–301. doi:10.1016/j.ehb.2011.03.004
- Nakamura, J., & Csikszentmihalyi, M. (2002). The concept of flow. In C. R. Snyder (Ed.), *Handbook of positive psychology* (pp. 89–105). New York, NY: Oxford University Press.
- Olshansky, S. J., Passaro, D. J., Hershov, R. C., Layden, J., Carnes, B. A., Brody, J., Hayflick, L., et al. (2005). A potential decline in life expectancy in the United States in the 21st century. *New England Journal of Medicine*, 352(11), 1138–1145. Retrieved from <http://www.nejm.org/doi/full/10.1056/nejmsr043743>
- Otlet, P. (1935). *Monde: Essai d'Universalisme*. Brussels: Mundaneum.
- Palmer, S. (2010). *Toxic Childhood: How The Modern World Is Damaging Our Children And What We Can Do About It*. Orion.
- Parunak, H. V. D. (2006). A survey of environments and mechanisms for human-human stigmergy. In D. Weyns, H. V. D. Parunak, & F. Michel (Eds.), *Environments for Multi-Agent Systems II* (pp. 163–186). Springer.
- Pinker, S. (2011). *The better angels of our nature: The decline of violence in history and its causes*. Penguin. Retrieved from <http://books.google.com/books?hl=en&lr=&id=c3cWagNsfMC&oi=fnd&pg=PT2&dq=pinker+angels&ots=aJJfnVMb7&sig=ZaRwXXnwNDqRKXZ1t9wSGM9M4oo>
- Proudfoot, D. (2012). Software Immortals: Science or Faith? In A. Eden, Johnny Søraker, Jim Moor, and Eric Steinhart (Ed.), *Singularity Hypotheses*, The Frontiers Collection (pp. 367–392). Springer. Retrieved from [http://link.springer.com/chapter/10.1007/978-3-642-32560-1\\_18](http://link.springer.com/chapter/10.1007/978-3-642-32560-1_18)
- Rodriguez, O. (2012). MOOCs and the AI-Stanford like Courses: two successful and distinct course formats for massive open online courses. *European Journal of Open, Distance, and E-Learning*. Retrieved from <http://www.eurodl.org/materials/contrib/2012/Rodriguez.htm>
- Sahlins, M. (2004). The original affluent society. *Investigating culture: an experiential introduction to anthropology*, 110.
- Schmidt, E., & Cohen, J. (2013). *The New Digital Age: Reshaping the Future of People, Nations and Business*. Random House Digital, Inc. Retrieved from [http://books.google.com/books?hl=en&lr=&id=Fl\\_LPoLdGKsC&oi=fnd&dq=The+New+Digital+Age+by+Eric+Schmidt+and+Jared+Cohen&ots=rly1gVnabZ&sig=mswnN7bo01oM9nh6J0ADEWRJBy8](http://books.google.com/books?hl=en&lr=&id=Fl_LPoLdGKsC&oi=fnd&dq=The+New+Digital+Age+by+Eric+Schmidt+and+Jared+Cohen&ots=rly1gVnabZ&sig=mswnN7bo01oM9nh6J0ADEWRJBy8)

- Steels, L., & Brooks, R. A. (1995). *The artificial life route to artificial intelligence: Building embodied, situated agents*. Lawrence Erlbaum.
- Stewart, J. (2000). *Evolution's Arrow: The Direction of Evolution and the Future of Humanity*. Chapman Pr.
- Stock, G. (1993). *Metaman: the merging of humans and machines into a global superorganism*. Simon & Schuster.
- Teilhard de Chardin, P. (1959). *The phenomenon of man*. Collins London.
- Thompson, C. (2011). How Khan Academy is changing the rules of education. *Wired Magazine*, (August), 126. Retrieved from [http://www.wired.com/magazine/2011/07/ff\\_khan/all/1](http://www.wired.com/magazine/2011/07/ff_khan/all/1)
- Thurner, S., & Poledna, S. (2013). DebtRank-transparency: Controlling systemic risk in financial networks. *Scientific reports*, 3. Retrieved from <http://www.nature.com/srep/2013/130528/srep01888/full/srep01888.html>
- Toffler, A. (1970). *Future shock*. New York: Random House.
- Turchin, V. (1995). A dialogue on metasystem transition. *World Futures*, 45, 5–57. Retrieved from <http://pespmc1.vub.ac.be/papers/Turchin/dialog.pdf>
- Turner, J. S. (2000). *The extended organism: the physiology of animal-built structures*. Harvard University Press. Retrieved from [http://books.google.be/books?hl=en&lr=&id=wapnxKbb23MC&oi=fnd&pg=PP3&dq=extended+organism&ots=QSfNw8kXZI&sig=q\\_0pZEwBHKa1tfxK90wDPJnqrb4](http://books.google.be/books?hl=en&lr=&id=wapnxKbb23MC&oi=fnd&pg=PP3&dq=extended+organism&ots=QSfNw8kXZI&sig=q_0pZEwBHKa1tfxK90wDPJnqrb4)
- Van Krevelen, D. W. F., & Poelman, R. (2010). A survey of augmented reality technologies, applications and limitations. *International Journal of Virtual Reality*, 9(2), 1. Retrieved from <http://kjcomps.6te.net/upload/paper1%20.pdf>
- Vidal, C. (2013). *The Beginning and the End: The Meaning of Life in a Cosmological Perspective*. Vrije Universiteit Brussel. Retrieved from arXiv preprint arXiv:1301.1648.
- Vinge, V. (1993). The coming technological singularity. *Whole Earth Review*, 88–95.
- Wierenga, E. (2012). Omniscience. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2012.). Retrieved from <http://plato.stanford.edu/archives/win2012/entries/omniscience/>
- Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *Evolutionary Computation, IEEE Transactions on*, 1(1), 67–82. Retrieved from [http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=585893](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=585893)
- Wright, R. (2001). *Nonzero: The logic of human destiny*. Vintage.
- Yanofsky, N. S. (2013). *The outer limits of reason: what science, mathematics, and logic cannot tell us*. MIT Press.
- Young, K. S. (1998). Internet addiction: The emergence of a new clinical disorder. *CyberPsychology & Behavior*, 1(3), 237–244. Retrieved from <http://online.liebertpub.com/doi/abs/10.1089/cpb.1998.1.237>
- Zander, T. O., Kothe, C., Jatzev, S., & Gaertner, M. (2010). Enhancing Human-Computer Interaction with Input from Active and Passive Brain-Computer Interfaces. In D. S. Tan & A. Nijholt (Eds.), *Brain-Computer Interfaces* (pp. 181–199). London: Springer London. Retrieved from [http://link.springer.com/10.1007/978-1-84996-272-8\\_11](http://link.springer.com/10.1007/978-1-84996-272-8_11)