

From Human Computation to the Global Brain: the self-organization of distributed intelligence

Francis Heylighen

Global Brain Institute
Vrije Universiteit Brussel

Introduction

The present chapter wishes to investigate the wider context of human computation, viewing it as merely one approach within the broad domain of distributed human-computer symbiosis. The multifarious developments in the “social” Internet have shown the great potential of large-scale collaborative systems that involve both people and information technologies. Here, I wish to explore this development in the broadest sense, as the self-organization of a distributed intelligence system at the planetary level—a phenomenon that has been called the “global brain”.

To get there, I will first define and illustrate the fundamental concept of distributed intelligence. Then I will review how such an intelligent network emerges and grows through a combination of self-organization and design. Finally, I will sketch some potential applications of the anticipated global brain.

Human-computer complementarity

The rationale for human computation is that people have certain intrinsic skills that are difficult to reproduce in computer programs. A computation system that makes use of those skills must therefore include people as information-processing agents. Thus, in human computation, people and computers are supposed to work together synergetically, the one complementing the other.

The reason for this complementarity lies in the fact that humans and computers process information in very different ways. Computers excel at accurately storing and retrieving discrete items, such as numbers or strings of characters. Human long-term memory, on the other hand is a network of associations that is continuously being modified by selective weakening, reinforcing, and combining of memory traces. As a result, people not only forget much of what they observed, but they are strongly biased in what they recall, and sometimes even “remember” things that never happened (Loftus & Pickrell, 1995). Thus, human memory is very unreliable compared to computer memory.

The problem gets even worse when people need to manipulate data, which happens in their *working memory*. This is the human equivalent of computer RAM. Working memory, however, cannot sustain more than some 4 items simultaneously (Cowan, 2001). Therefore, most people are unable to make any but the most trivial calculations without the help of pen and paper. Computers, on the other hand, are virtually unlimited in the amount of items they can manipulate, and do not make mistakes when retrieving stored information.

This unreliability of human memory is compensated by the fact that the neural networks that make up our brain are very effective at learning asso-

ciations between different experiences, and thus uncovering subtle patterns in information (McLeod, Plunkett, & Rolls, 1998). Moreover, the brain is remarkably reliable in the *recognition* of patterns similar to patterns experienced before, even while being poor at *recall*, i.e. retrieving exact data. Recognition is so robust because the newly activated pattern can be very different from the one stored in memory, but still the activation spreads through a myriad of learned associations until it activates memories that are in some way related to the new perception. This prepares the mind to anticipate features of that pattern analogous to features experienced before, a form of “intuition” that computers generally lack.

Moreover, human cognition is *situated* and *embodied* (Anderson, 2003; Clancey, 1997; Clark, 1998): we continuously interact with our environment via exquisitely sensitive and sophisticated sensory organs and muscle systems, which have evolved over billions of years. This provides our brain with a very high-bandwidth channel for input, output and feedback, allowing it to learn the high-dimensional, fine-grained patterns and correlations that characterize the real world with all its complexities and dynamics. Thanks to this direct coupling between brain and outside world we learn not only to recognize subtle patterns, but to perform precisely coordinated actions. Indeed, the fine-grained sensory feedback we get allows us to automatically perform the kind of complex manipulations that are so difficult for robotic devices.

This on-going interaction has provided people with a lifetime of real-world experience, getting them to know subtle relations between millions of phenomena, variables and stimuli. The resulting knowledge is nearly impossible to implement in a computer program, as most of it is too fuzzy, holistic and context-dependent to be exteriorized in the form of symbols and rules. The difficulty of formalizing such knowledge is known in AI as the

“knowledge acquisition bottleneck” (Wagner, 2006). It is one of the reasons that information technologists have turned to systems that include *human computation*: letting people perform those tasks that are too difficult for a computer program, while using computers to do the tasks that are difficult or tedious for people (Quinn & Bederson, 2011; Von Ahn, 2009).

Distributed intelligence

Human computation is one among a variety of paradigms that study how people supported by information and communication technologies are able to solve more problems together than when working in isolation. These approaches have been variously called man-machine interaction, human-computer symbiosis, computer-supported cooperative work, social computing, crowdsourcing, collective intelligence, and wisdom of crowds. Different labels or metaphors tend to emphasize different aspects of this synergistic interaction, while ignoring others. For example, human computation considers individuals as merely components performing specific subroutines within a clearly defined overall program, thus making them subordinate to the technological system (Nagar, 2011). Computer-supported cooperative work, on the other hand, takes the opposite stance, seeing the technology as subordinate to the human interaction, while collective intelligence a priori ignores technology, even though its practical implementations almost always rely on some kind of information technology.

My aim here is to look for the most general collaborative framework, implying a minimal bias about what kind of activity is performed by whom or by what. A good starting point can be found in the concepts of *information processing* and *distributed cognition* (Heylighen, 2012a; Nagar,

2011). Information processing is a notion broader than computation, as it does not assume well-structured, “computable” tasks. It has been extensively used to analyze how organizations solve their problems (Galbraith, 1974; Tushman & Nadler, 1978)—with or without computer technology. Neural network models illustrate how information processing in the brain is *distributed* (Rumelhart & McClelland, 1986): different neurons deal simultaneously with different aspects of the information, while aggregating their results into a comprehensive interpretation. Such collaboration between different components of the brain inspired Minsky (1988) to conceive of the mind as a “society” of interacting agents. The analogy works both ways, though: human society itself is in a number of respects similar to a brain, since it consists of agents that together solve problems that are too difficult for the agents individually. Thus, the distributed information processing perspective is applicable at all levels, from neural circuits, via the brain and organizations, to society as a whole. The principle is simply that a collective of collaborating agents can process more complex information more extensively than any individual member of that collective.

The last step we need to reach the notion of distributed cognition is to observe that physical objects or tools too can function as information processing agents. The simplest tools, such as books, merely store information for later use, thus compensating for the unreliability of human memory. Other tools, such as telephones, can transfer information from one agent to another across distances. Yet other tools, such as sensors, cameras or recorders, can capture external information. The most sophisticated tools—as exemplified by modern ICT—store, register, transfer and process information. Integrate such tools with human agents into a coordinated system or organization and the result is *distributed cognition* (see Fig. 1): the acquisition, propagation and processing of information across a heteroge-




Figure 1: a depiction of an organization as a distributed cognitive system, i.e. a network of humans and artifacts that store, process, and propagate information along the links in the network. The thickness of an arrow represents the intensity of the communication across the corresponding link. Incoming arrows represent input to the system (its perception of the environment), outgoing arrows its output (its action on the environment).

neous network of people and artifacts (Dror & Harnad, 2008; Hollan, Hutchins, & Kirsh, 2000; Hutchins, 2000).

Distributed cognition as originally conceived by Hutchins (1995, 2000) is a description of an existing situation: social systems have always used external aids for propagating and processing information. What the newer approaches, such as human computation, aim at is to use information technologies to make such distributed processing much more powerful, more efficient, or more intelligent. Let us then call this new endeavor *distributed intelligence* (Fischer, 2006; Heylighen, 2012a). The present paper wishes

to investigate the future of distributed intelligence: how are distributed intelligence technologies likely to develop and to affect society at large? To answer that question, we must first understand how distributed intelligence emerges from its components.

Self-organization

Distributed intelligence can be understood as the coordinated activity of a collective of agents (human or technological) that process and propagate information between them. In formal organizations, such as firms, computer systems, or administrations, such coordination is normally the result of *design* (Galbraith, 1974; Tushman & Nadler, 1978). This means that some person or group of persons has developed a scheme that specifies which information is to be processed by which agent, and how the output of that process is then sent for further processing to one or more other agents. Such schemes take the form of computer programs, organizational charts or workflow diagrams.

However, as anybody who has worked in an organization knows, such a scheme only captures a small part of the actual information flow. Most communication follows informal channels, which together form a social network. A social network is formed by links of acquaintanceship, friendship or trust, which are built up through the personal encounters and experiences of the people in the group. In other words, a social network is not imposed by central design, but emerges through decentralized *self-organization*. If we zoom out and consider increasingly large distributed cognitive systems, we will notice that explicit organization plays an increasingly small role, while implicit networks become increasingly more important. The reason is simply that the more complex the system, the

more difficult it becomes to completely specify the rules about which component is to work with which other component in which way. If we compare the poor results of central planning in communist societies with the effectiveness of the “invisible hand” of the market, then we can only conclude that self-organization must be the major driver of coordination in a system as complex as society.

Self-organization is not just the foundation on which social systems are built. Its power is increasingly being harnessed for building technological systems. Here too, designers are confronted with a complexity bottleneck: as soon as the number of components and their interactions become too large and/or too variable, explicit design or “programming” of the system becomes infeasible. That is why computer scientists and engineers are now exploring self-organizing solutions to the problem of how to coordinate a variety of interacting modules (Bartholdi III, Eisenstein, & Lim, 2010; Dressler, 2008; Elmenreich, D’Souza, Bettstetter, & de Meer, 2009).

Self-organization is perhaps most critical in the Internet, which is the most complex socio-technological system that presently exists. It is simply impossible to make a rational design for how the different websites and services on the Internet should be connected, because no one knows exactly which services exist and what they can do. Moreover, thousands of new pages, forums and applications appear every day, seeking their place within an anarchic and highly competitive network of linked information sources. Thus, the topology of cyberspace is changing so rapidly that no central authority can ever hope to control it.

How does self-organization work? At the most basic level, every evolutionary process uses *trial-and-error*, a mechanism that can be described more accurately as *blind-variation-and-selective-retention*. If you do not

know how to fit things together, then you try a variety of combinations. You then eliminate the ones that do not work (errors), and select the others for retention. This process is iterated: the retained solutions are again modified, producing some variants that work better and are therefore retained, some that work worse and are therefore rejected. If you continue this iteration long enough, you are bound to end up with something much better than what you started out with. This process can be speeded up with the help of positive feedback: amplifying or multiplying the “good” solutions in proportion to their “goodness”, so as to increase the average quality of your starting material for the next iteration, but without losing the necessary variety. This is the mechanism underlying both biological evolution and its application to computation as implemented e.g. by genetic algorithms (Booker, Goldberg, & Holland, 1989).

The same kind of positive-feedback enhanced iteration occurs in self-organizing systems (Heylighen, 2013), with the difference that the selection is “natural”. That means that there is not any agent or explicit criterion deciding what to keep and what to reject. It is rather the system as a whole that determines what survives and what is eliminated. The selected variations are the ones that *fit* their environment. In the complex adaptive system that we are considering, the environment of a component is constituted by the agents that it directly interacts with. A “fit” interaction in this case means one that is beneficial, so that the participating agents are inclined to continue it. If the interaction is not beneficial, then there is no reason to maintain it, and the link with the component will be cut. Thus, natural selection here is in the first place a selection of links between components or nodes in the network. The same component may fit in well with certain agents, but not with others. To find out where it fits best, it needs to try out various links, keeping (or strengthening) the good ones and eliminating (or

weakening) the less good ones. This is the same mechanism that underlies learning in the brain: useful links (as embodied by synapses connecting neurons) are reinforced; less useful ones are weakened, and eventually cut.

The self-organization of distributed intelligence

Let us now apply this self-organizing dynamics to heterogeneous networks of cognitive agents, i.e. people and ICT systems. Human computation systems are examples of such heterogeneous networks, albeit that their organization is largely designed or programmed. At the level of the Internet as a whole, however, size and heterogeneity increase to such a degree that design must make place for self-organization via *selective linking*. A simple illustration of how this can happen is bookmarking: when a person surfing the web encounters a particularly interesting or useful page, such as a weather forecasting service, a search engine, or an overview of the domain in which the person is interested, then that person will store a link to that page in the browser, as a “favorite” or “bookmark”. This makes it easy for the person to come back frequently to that page. Here, a stable link is created between a human and an ICT agent.

A link between two human agents is created when one person meets another one—face-to-face or on the web—and finds that person interesting enough to add him to her list of “contacts” in some social network application, such as Facebook or LinkedIn. This link now makes it easy for the first person to directly pass on information to the second one. A connection between two ICT agents is established when a hyperlink is made from one webpage or website to another one, or when one ICT system (say, the

Facebook platform) decides to exchange data with another one (say, the Skype calling service).

In all these cases, links that are successful, in the sense that the agents benefit from them, will survive and be reinforced, while links that are useless or counterproductive will be forgotten and eventually erased. For example, your link to site A may turn out to be particularly useful, and therefore you give it a more prominent place, making the one to the less user-friendly site B redundant, so that you eventually remove it. Similarly, you may from time to time remove “contacts” that turn out to be tiresome, while upgrading others to the status of “friends” or “collaborators”. This on-going variation and selection of links makes the network as a whole evolve towards an increasingly efficient or “intelligent” organization. This is analogous to the way the neural networks in our brain learn how to respond more intelligently to the problems they encounter.

The intelligence of this distributed system can be understood through the paradigm of *challenge propagation* (Heylighen, 2012a). A problem, question, message or opportunity constitutes a *challenge* for one or more agents: it incites the agent to act, i.e. to respond in a way that may solve the problem, answer the query, reply to the message, or seize the opportunity. A challenge in this sense is a generic term for a piece of information that carries value for an agent, and that therefore can motivate the agent to process the information in order to extract that value. Challenges can be positive (acting on them provides benefit: *opportunities*) and/or negative (not acting on them makes you lose benefit: *problems*). Dealing with challenges is therefore a generalization of solving problems. Problem solving is considered to be the hallmark of intelligence, as more intelligent agents can solve more or more difficult problems than less intelligent ones.

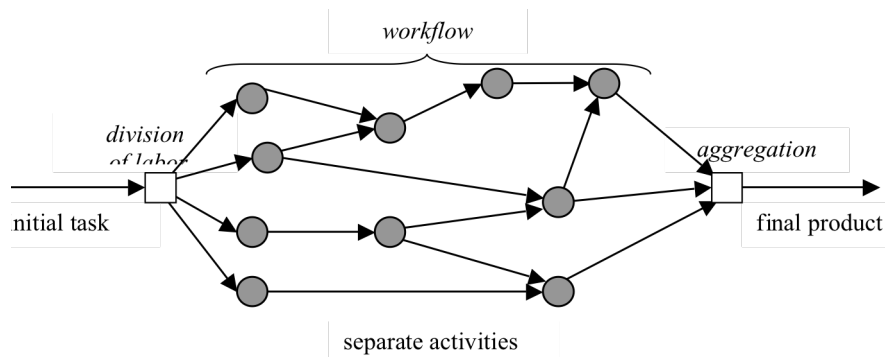


Figure 2: An illustration of coordination, in which an initial task is split up in separate activities performed by different agents (division of labor), which are followed by other activities (workflow), and whose results are assembled into a final product (aggregation). Grey circles represent individual agents performing activities. Arrows represent the “flow” of challenges from one agent to the next.

To measure the intelligence of a distributed network, we can then try to establish its capacity to effectively process challenges. Normally, different agents have different skills in dealing with challenges. For example, computers excel in making complex calculations, while people excel in understanding spoken language. Different people and different computer agents have further their own special abilities, so that our network as a whole will present a wide range of finely grained skills and expertise. A complex challenge (say, global warming) has a large number of aspects that each require different skills. The problem now is to *distribute* the different challenge components across the different agents so as to make sure the challenge as a whole is dealt with in an efficient way. This is the basic problem of coordination, which includes *division of labor* (who deals with what challenge component?), *workflow* (where does a component go after it has been partially dealt with?), and *aggregation* (how are all the finished

pieces of work assembled?) (Heylighen, Kostov, & Kiemen, 2013; Heylighen, 2013) (see Fig. 2).

Perhaps surprisingly, such distributed coordination can self-organize relatively easily across the Internet, via the mechanisms of stigmergy (Heylighen et al., 2013; Heylighen, 2007a) and challenge propagation (Heylighen, 2012a). A good illustration can be found in the different open source communities developing complex software without central authority, and Wikipedia, the online encyclopedia created and maintained by millions of volunteer contributors. *Stigmergy* is an implicit coordination mechanism whereby a challenge left by an agent in a workspace that is shared with other agents stimulates those agents to continue dealing with that challenge (Parunak, 2006). For example, a paragraph added to a Wikipedia page by one person may incite a second person to add some extra details, a third one to add a reference for the new material, and a fourth one to edit the text so as to make it more readable. The reference may then be checked and correctly formatted by a software agent. In this case, challenges are spontaneously addressed by subsequent agents mediated by the shared workspace (in this case the Wikipedia website). In the case of *challenge propagation*, the workflow is initiated by the agents themselves. An example is an email message sent and forwarded with comments by different people, a “post” in a social network or forum that is reposted to other forums, or a task that is proposed by a crowdsourcing system to people interested to work on it.

In both cases, challenges can travel more or less efficiently across the network of agents and workspaces until they find an agent able and willing to deal with them, and then continue their journey along other agents dealing with the remaining aspects. This allows complex challenges to be resolved in a distributed manner, by harnessing the collective intelligence of the dif-

ferent components (human and technological) of the network. Presently, my research group is developing a mathematical model of this process, in order to investigate precisely how the distributed intelligence of the network can increase as it selectively strengthens or weakens its links (Heylighen, Busseniers, Veitas, Vidal, & Weinbaum, 2012).

The Global Brain

What happens when such a self-organizing distributed intelligence network grows to encompass the planet (as the Internet already does)? The result can perhaps best be understood using the metaphor of a *global brain* (Bernstein, Klein, & Malone, 2012; Goertzel, 2002; Heylighen, 2008; Mayer-Kress & Barczys, 1995). The global brain can be seen as the nervous system of the *planetary superorganism* (De Rosnay, 2000; Heylighen, 2007b; Stock, 1993). This is the “living system” (Miller, 1995) formed by all people on this planet together with their artifacts and technologies. The task of its brain is to gather and process information about the situation of the world and all its people, find solutions to any problems it detects, and incite and coordinate actions to deal with those challenges (cf. Helbing, Bishop, Conte, Lukowicz, & McCarthy, 2012). This is similar to the task of the human brain, which gathers information through its sensory organs, processes that information in order to evaluate the situation, then reflects about strategies to deal with the challenges it finds, and finally implements those strategies by sending signals to the muscles so as to direct and coordinate their actions. A secondary task of both human and global brain is to learn from its experiences by reinforcing the successful links in its network (and weakening the others). This allows it to develop ever more detailed

and accurate knowledge about itself and the environment in which it lives, and thus to become ever better at dealing with the challenges it encounters.

We should expect the problem-solving abilities of the global brain to be orders of magnitudes larger than that of any single individual, organization, or computer system. This is because all people and computers collectively have access to immensely more knowledge and processing capacities than any of them individually (Heylighen, 2012b). The only requirement to efficiently harness this collective intelligence is coordination. This can be expected to self-organize, as illustrated by both empirical observations (Heylighen, 2013; Woolley, Chabris, Pentland, Hashmi, & Malone, 2010) and simulations (Elmenreich et al., 2009; Heylighen et al., 2012). However, self-organization at a scale as large as the world obviously needs time, as countless iterations of the variation and selective reinforcement process must take place, and as any provisionally “fit” result will need to be updated as soon as a new agent or technology appears on the scene. Thus, all components of the global network will continue to co-evolve at a rapid pace, increasing their degree of coordination, efficiency and intelligence in the process, but in a manner so complex that we cannot predict it in any detail.

It is impossible to say when this process will have produced the equivalent of a global brain, since distributed intelligence is a continuously growing and evolving measure of coordination, not a phenomenon that either is or is not present. Thus, we cannot “detect” the presence or absence of a global brain, but we can conceivably measure the increase in distributed intelligence of the global network. In our mathematical model (Heylighen et al., 2012), we have developed one such quantitative measure, and suggested some methods to gather the necessary empirical data to test its evolution in the real world—but these are very preliminary results.

The self-organization of the global brain could in principle be accelerated by complementing it with thoughtful design. As we start to better understand processes such as self-organization, distributed cognition, collective intelligence and human-computer complementarity, we may be able to avoid some of the trial-and-error search, and develop systems that produce coordinated information processing more quickly and more reliably. For example, inspired by their insights into collective intelligence, Bernstein et al. (2012) have suggested methods for “programming” the global brain—but these too are very preliminary. Further methods may be discovered through research in human computation, crowdsourcing, ontology development, and related fields. However, no single system, method or program will ever be able to capture the immense size and complexity of our planetary network. Therefore, we must resign ourselves to the fact that we will never be able to fully control the process. Perhaps the most promising approach is what may be called “guided self-organization” (Helbing, 2012): developing strategies, programs, and institutions that will facilitate and to some degree steer the self-organization of the global brain in what appear to be the most fruitful directions. But to achieve that, we must first of all better understand what the global brain would be able to do, and especially what we want it to do.

Some applications of the Global Brain

Now that we have a better grasp of how a global brain-like system would emerge, let us try to sketch some of its potential benefits for society. In principle, the Global Brain should help us to tackle any individual or collective challenge, by providing us with a vast reservoir of knowledge, sen-

sory data, information processing capacity, and ability to incite coordinated action. A first domain that would profit from these superhuman abilities is the economy.

The market is the collective system of transactions that helps supply to match demand, and thus to satisfy the public's need for products and services. A traditional market is rather inefficient, requiring a huge infrastructure of middlemen, specialized organizations such as stock exchanges and auctions, and communication channels. The Internet already allows such transactions to take place much more quickly and transparently, with less cost and effort. This strongly reduces friction, making the economy more efficient so that demand can be satisfied more rapidly, more accurately, and at a lower cost. The global brain will not only facilitate communication between suppliers and clients, but help buyers to find the best value (e.g. through shopping agents to find items and compare prices), and help sellers to get the best price (e.g. through auctioning systems and targeted advertisements).

The net effect will be that growth and productivity increases, while inflation and economic instability decrease. Moreover, there will be less waste because of unsold items or goods shipped far away when there is demand around the corner. The direct incorporation of collective effects ("externalities") in the decision-making process will moreover allow a more efficient governance over the economy, thus protecting employees and consumers while reducing inequalities and pollution, without the added complexity, bureaucracy and rigidity that tend to accompany such interventions in a centralized political system.

The global brain can moreover help eliminate conflicts. It in principle provides a universal channel through which people from all countries, lan-

languages and cultures of this world can converse. This makes it easier to reduce mutual ignorance and misunderstandings, or discuss and resolve differences of opinion. The greater ease with which good ideas can spread over the whole planet and the collective improvement on those ideas will make it easier to reach global consensus about issues that concern everybody. The free flow of information will make it more difficult for authoritarian regimes to plan suppression or war. The growing interdependence will stimulate collaboration, while making war more difficult. The more efficient economy will indirectly reduce the threat of conflict, since there will be less competition for scarce resources.

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

Similarly positive effects can be conceived in domains as diverse as health, well-being, democratic participation, sustainable development, work productivity, disaster prevention and relief, education, research, innovation, industry, traffic, logistics, and ecosystem management (Heylighen et al., 2013; Heylighen, 2002, 2007b, 2007b). There seems to be no end to the potential applications of a distributed intelligence system at the world level. Many of these applications are already becoming apparent in the present Internet, but their beneficial effect is held back by the general con-

fusion, information overload and uncertainty that accompanies the present explosion in new technologies and functions (Heylighen et al., 2013). It is to be expected that the overall benefits will multiply as the network becomes more streamlined and intelligent, and the agents using it more coordinated in their activities. Then, only the sky will be the limit to what a global brain can achieve...

References

- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial intelligence, 149*(1), 91–130.
- Bartholdi III, J. J., Eisenstein, D. D., & Lim, Y. F. (2010). Self-organizing logistics systems. *Annual Reviews in Control, 34*(1), 111–117. doi:10.1016/j.arcontrol.2010.02.006
- Bernstein, A., Klein, M., & Malone, T. W. (2012). Programming the Global Brain. *Communications of the ACM, 55*(5), 1.
- Booker, L. B., Goldberg, D. E., & Holland, J. H. (1989). Classifier systems and genetic algorithms. *Artificial intelligence, 40*(1-3), 235–282.
- Clancey, W. J. (1997). *Situated cognition: On human knowledge and computer representations*. Cambridge Univ Pr.
- Clark, A. (1998). Embodied, situated, and distributed cognition. *A companion to cognitive science, 506–517*.

- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(01), 87–114.
- 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>
- Dressler, F. (2008). A study of self-organization mechanisms in ad hoc and sensor networks. *Computer Communications*, 31(13), 3018–3029.
- Dror, I. E., & Harnad, S. R. (2008). *Cognition Distributed: How Cognitive Technology Extends Our Minds*. John Benjamins Publishing.
- Elmenreich, W., D'Souza, R., Bettstetter, C., & De Meer, H. (2009). A Survey of Models and Design Methods for Self-Organizing Networked Systems. *Self-Organizing Systems*, 37–49.
- Fischer, G. (2006). Distributed intelligence: extending the power of the unaided, individual human mind. In *Proceedings of the working conference on Advanced visual interfaces* (pp. 7–14). New York, NY, USA: ACM. doi:10.1145/1133265.1133268
- Galbraith, J. R. (1974). Organization design: An information processing view. *Interfaces*, 4(3), 28–36.
- Goertzel, B. (2002). *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers.

- Helbing, D. (2012). Managing Complexity. In D. Helbing (Ed.), *Social Self-Organization* (pp. 285–299). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-24004-1_15
- 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.
- Heylighen, F. (2002). The global brain as a new utopia. In *Zukunftsfiguren*. Suhrkamp, Frankfurt. Retrieved from <http://pespmc1.vub.ac.be/papers/GB-Utopia.pdf>
- Heylighen, F. (2007a). 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. (2007b). The Global Superorganism: an evolutionary-cybernetic model of the emerging network society. *Social Evolution & History*, 6(1), 58–119.
- Heylighen, F. (2008). Accelerating socio-technological evolution: from ephemeralization and stigmergy to the global brain. In *Globalization as evolutionary process: modeling global change* (p. 284). Routledge.

- Heylighen, F. (2012a). *Challenge Propagation: a new paradigm for modeling distributed intelligence* (No. 2012-01). Brussels, Belgium. Retrieved from <http://pcp.vub.ac.be/papers/ChallengePropagation.pdf>
- Heylighen, F. (2012b). 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.
- Heylighen, F. (2013). Self-organization in Communicating Groups: the emergence of coordination, shared references and collective intelligence. In À. Masip-Bonet & A. Bastardas-Boada (Eds.), *Complexity Perspectives on Language, Communication and Society* (pp. 117–149). Berlin, Germany: Springer. Retrieved from <http://pcp.vub.ac.be/Papers/Barcelona-LanguageSO.pdf>
- 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). Retrieved from <http://pespmc1.vub.ac.be/papers/TowardsGB-model.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.), *Routledge Handbook on Knowledge Economy, Education and Digital Futures*. Routledge. Retrieved from <http://pcp.vub.ac.be/Papers/MobilizationSystems.pdf>

- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(2), 174–196.
- Hutchins, E. (1995). *Cognition in the Wild* (Vol. 262082314). Cambridge, MA: MIT press. Retrieved from http://books.google.be/books?id=AfupQgAACAAJ&dq=Cognition+in+the+Wild&hl=en&sa=X&ei=YjLWT6PGLaOn0QWwyqiBBA&redir_esc=y
- Hutchins, E. (2000). Distributed cognition. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences*. Elsevier Science.
- Loftus, E. F., & Pickrell, J. E. (1995). The formation of false memories. *Psychiatric Annals*, 25(12), 720–725.
- 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.
- McLeod, P., Plunkett, K., & Rolls, E. T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford University Press Oxford.
- Miller, J. G. (1995). *Living systems*. University Press of Colorado.
- Minsky, M. (1988). *The society of mind*. Simon & Schuster.
- Nagar, Y. (2011). Beyond the Human-Computation Metaphor. In *Privacy, security, risk and trust (passat), 2011 ieee third international conference on*

and 2011 ieee third international conference on social computing (socialcom) (pp. 800–805). Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6113219

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.

Quinn, A. J., & Bederson, B. B. (2011). Human computation: a survey and taxonomy of a growing field. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (pp. 1403–1412). Retrieved from <http://dl.acm.org/citation.cfm?id=1979148>

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. San Diego: University of California Press.

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

Tushman, M. L., & Nadler, D. A. (1978). Information processing as an integrating concept in organizational design. *Academy of Management Review*, 613–624.

Von Ahn, L. (2009). Human computation. In *Design Automation Conference, 2009. DAC'09. 46th ACM/IEEE* (pp. 418–419). Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5227025

Wagner, C. (2006). Breaking the knowledge acquisition bottleneck through conversational knowledge management. *Information Resources Management Journal (IRMJ)*, 19(1), 70–83.

Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *science*, 330(6004), 686–688.