

# **Distributed Intelligence Technologies: present and future applications**

Francis Heylighen

*Global Brain Institute,*

*Vrije Universiteit Brussel*

This paper surveys a wide variety of potential applications of *distributed intelligence technologies*. These are systems that tackle complex challenges in a distributed manner, by collecting, processing and routing information and actions that are spread across a global network of human and technological agents. The eventual integration of all these systems is expected to produce a distributed intelligence at the planetary scale, the *Global Brain*. Supporting technologies for such distributed problem solving include Wikipedia, the Semantic Web, the Internet of Things and the Smart Grid. These and future extensions will facilitate the propagation and coordination of information, knowledge, energy, physical objects, actions and personal identity. By extrapolating from existing applications such as social media, data mining, online shopping and MOOCs, it is argued that these technologies can in principle satisfy all our material, health, safety, social, achievement, cognitive, and even self-actualization needs.

## **1. Introduction**

Over the past two decades, the Internet has taken over ever more social, economic and technological functions from other communication systems, and this at an absolutely staggering speed. At the same time, it has been opening up a seemingly infinite variety of new applications. People use the Internet for applications as diverse as ordering groceries, organizing political rallies, watching movies, financing new inventions, selling second-hand goods, discussing global problems, keeping in touch with family, monitoring buildings remotely, guiding self-driving cars, publishing photos, articles and books, keeping stock in warehouses, distributing intricate calculations across thousands of independent computers, tracking public transport, exchanging scientific information,

“crowdsourcing” tasks to anonymous workers, and remotely following courses at prestigious universities. Yet, for every successful new service or application, plenty of equally promising new ventures seem to fail.

This explosion in the number of actual and potential applications of the Internet is overwhelming. The resulting confusion makes it very difficult to discern stable trends—except for a general growth in Internet use. Long-term prediction of these multifarious developments seems especially daunting. Yet, there exists a paradigm that promises to bring some order to this tangle of seemingly chaotic, yet interdependent developments: the *Global Brain* (Bernstein, Klein, & Malone, 2012; Goertzel, 2002; Heylighen, 2011; Mayer-Kress & Barczys, 1995; Russell, 1995).

The Global Brain paradigm conceives the Internet as the nervous system of a *global superorganism*, a planetary system consisting of all humans, their artifacts, and the social, economic and technological links that tie them together (De Rosnay, 2000; Heylighen, 2007a; Stock, 1993). The idea is that global exchanges have made the people on this planet interdependent to such a degree that together they form a single “living system” (J. G. Miller, 1965, 1995). The function of the nervous system is to coordinate the different activities taking place in the organism. More in particular, a brain needs to gather information about what is going on in and around the system, process the information in order to decide what this means and what should be done about it, and finally initiate appropriate actions to deal with the perceived challenges.

Contemporary science sees societies, organisms and brains as *complex adaptive systems* (Ball, 2012; Holland, 1992; J. H. Miller & Page, 2007). This means that they consist of a vast number of relatively autonomous agents (such as cells, neurons or individuals) that interact locally via a variety of channels. Together these channels form a complex, dynamic network. Out of these non-linear interactions, some form of coherent, coordinated activity emerges—a phenomenon known as *self-organization* (Camazine et al., 2003; Heylighen, 2013). The resulting organization is truly *distributed* over the components of the system: it is not localized, centralized or directed by one or a few agents, but arises out of the interconnections between all the agents.

In the case of the Global Brain (GB), the self-organization of this distributed capacity for information processing is still in the early stages: the global system is far from having settled into a stable, coordinated regime. However, the breakneck speed that characterizes the spread of the Internet clearly exemplifies the accelerating pace of such a process of global coordination (Heylighen, 2008, 2016), which is driven by the positive feedback of new applications enabling and inspiring further applications. While such a non-linear process cannot be predicted in any detail, the Global Brain paradigm provides us with a long-term, qualitative view of the most likely outcome: an integrated, intelligent system for processing all the information relevant for the survival and development of the planetary superorganism (Heylighen, 2015a).

The present paper will start from this broad view in order to derive a more concrete list of present and future applications of the Internet. The emphasis will be on the distributed character of global information processing, because this is what most fundamentally distinguishes the new paradigm from the older paradigm, which sees political, economic or cognitive organizations as centralized, hierarchical systems. Therefore, we will first investigate what it means for an information-processing system to be distributed. Then, we will examine in more detail some of the specific technologies that will be needed to support a distributed intelligence at the global level. Finally, we will survey how such distributed intelligence can satisfy the basic needs of human individuals. This will give us a more concrete view of some major applications that can be foreseen.

## **2. Distributed intelligence**

### **2.1. *From individual to distributed intelligence***

To better understand which kind of applications would be part of a global brain, we need to define *distributed intelligence* (Fischer, 2006; Heylighen, 2014a). This will allow us to distinguish technologies that are effectively founded on some form of distributed intelligence from those

that merely support traditional functions. For example, a centralized database to which people can get individual access via the Internet to check records is not essentially different from the card catalogs that were used in libraries before the advent of computers. Similarly, using your smartphone to send a photo to a friend is not essentially different from sending a postcard via the mail. These Internet applications are merely more efficient versions of traditional communication media, and will simply continue to become quicker, cheaper and easier to use.

We will define *intelligence* as “the ability to gather and process information so as to efficiently solve problems and exploit opportunities”. What are considered problems, opportunities—or more generally *challenges* (Heylighen, 2014a)—will depend on the goals and values of the decision-maker, who can be an individual organism, an organization, or the global superorganism. Efficiently dealing with a challenge means selecting and performing the right actions that solve the problem or exploit the opportunity. It is the process of adequate selection that is the essence of intelligence (Ashby, 1956).

Traditional models of intelligence in cognitive science and artificial intelligence see the process of problem solving as a sequential search through a space of potential solutions. The attempts to simulate the neural networks used by our brain, however, led to the notion of parallel, *distributed* processing of information (Bechtel & Abrahamsen, 1991; McLeod, Plunkett, & Rolls, 1998; Rumelhart & McClelland, 1986). The idea is that different units or “neurons” deal simultaneously with different aspects of the problem or question. In other words, the problem is split up into parts or aspects that are processed by several autonomous agents (active units) working in parallel, without central supervision or direction. Their contributions are then reassembled or aggregated into a collective solution.

A fundamental advantage of this approach is flexibility and robustness. The many contributions ensure redundancy of function: individual units may be unavailable, produce erroneous results, or lack relevant data, but the resulting errors tend to be compensated by the signals coming from the other units, so that the aggregate result normally is informative—even in the most confused situations. In a centralized, sequential process, on the other hand, a single malfunction along the line

can be sufficient to throw everything off-course, so that no useful result is produced. Another advantage is that the different agents can represent a wide variety of skills, perspectives, or experiences, thus allowing a balanced, integrated approach of the most complex problems.

The same mechanism of compensating for individual ignorance or bias by aggregating a large variety of contributions characterizes successful applications of *collective intelligence* (Heylighen, 1999, 2013; Malone, Laubacher, & Dellarocas, 2010; Surowiecki, 2005). But in typical social systems, distributed intelligence is more than collective intelligence: contributions do not only come from the people in a collective, but from a variety of artifacts, tools and technologies that sense, register, store, process or transfer information. This is the perspective of *distributed cognition*, originally proposed by the ethnographer Hutchins (Clark, 1998; Heylighen, Heath, & Van Overwalle, 2004; Hutchins, 2000). In real-world problem solving, we routinely rely on tools such as pen and paper, maps, cameras, telephones and calculators to gather and process information. We also rely on other people to provide us with their unique observations, skills or ideas. For a complex system—such as a Navy ship (Hutchins & Lintern, 1995)—to function well, all the people and artifacts involved need to work together in a *coordinated* manner, by sending the right messages at the right moments to the right destinations.

## **2.2. Emergence of distributed intelligence**

In truly complex systems, such distributed coordination is typically the result of *self-organization* (Heylighen, 2013), not of central planning. However, self-organization is usually a slow and difficult process that needs to overcome a variety of obstacles, given that it needs to produce global coordination out of local interactions between agents that have only a very limited perspective on the whole. This is where Internet technologies can play a crucial role. They make it in principle possible for any agent (human or artificial) to interact in real-time with any other agent, while keeping a detailed trace of such interactions and their outcomes. This makes it easier to find, select and reinforce the

interactions that are most effective in tackling problems, while eliminating the less useful ones (Heylighen, 2008; Heylighen, Busseniers, Veitas, Vidal, & Weinbaum, 2012). As a result, coordination and distributed intelligence can evolve much more widely, quickly and effectively.

Internet technology not only facilitates globally distributed intelligence, it makes it virtually inevitable. Suppose that there exist two systems that perform at first sight separate functions, independently of each other. Imagine, for example, a medical database keeping track of individuals' disease records, and a cellular phone network, routing calls from one person to another. The Internet makes it possible to connect them together, but why would anybody want to do that? A fundamental reason for interconnection is that the people who use these technologies are themselves interconnected, and so are their problems—simply because they are all members of the same planet-wide society. There is not a single aspect of one person's life that is not potentially relevant for another aspect of another person's life. For example, my medical record may be relevant to your cellular phone record, because I carry a highly contagious disease and your cellular phone use shows that you were in the same place as I was at the same time. Or perhaps the record of our communications shows that we have both closely interacted with the same friend, who may well be the source of the infection.

Imagine that a system is developed to interconnect the two first systems—say, the medical database and the cellular network. Initially, this interconnection may exhibit some shortcomings, like poor protection of privacy, or vulnerability to computer viruses or to hacking. Through trial-and-error (which is the same as variation and natural selection), flawed versions will eventually be replaced by better versions. At that moment, we have an integrated system that is potentially much more powerful than the two initially disconnected systems. In the present example, the interconnection makes it possible to track the spread of infectious agents, and to warn potential carriers of their risk of infection.

The same reasoning extends to all systems that are not interconnected yet: sooner or later an interface will be developed between them, which allows the one system to benefit from the information produced by the other, and vice versa. The principle is simple: because

all individuals and systems are part of a coherent global superorganism, every piece of information produced in one part of that superorganism is potentially useful when decisions need to be made in some other part. Without interconnection, this information would remain inaccessible to the other components of the global system. As a result, these other parts would lack some relevant data, and therefore produce suboptimal decisions.

This principle explains why links between systems and components are more likely to be added and reinforced than to be removed. Eventually, all systems will become interconnected, and their pattern of interconnection will become increasingly efficient, maximizing synergies while minimizing conflicts or frictions (Heylighen, 2008, 2013, 2016), so that information can take the most direct route from those that produce it to those that need it. That is why there is a continuing trend for all information technologies to merge into one giant distributed intelligence.

### ***2.3. Aspects of distributed intelligence***

Let us analyze more concretely what is needed for a technology to support distributed intelligence. Intelligent information processing can be decomposed into the following stages:

1. **input**: collecting a variety of data about phenomena that are potentially relevant for the system (this is the sense of “intelligence” as information gathering)
2. **processing**: aggregating, filtering and recombining the input information so as to recognize the challenges most important to the system, and developing strategies to deal with those challenges (this is the sense of “intelligence” as interpretation and problem solving)
3. **output**: sending the derived decisions to the components of the system that are in contact with the outside world, so that they can address the observed challenges.

4. **feedback**: monitoring the effect of the performed actions, and using this information as new input for a next iteration of the process, so as to correct for mistakes and unforeseen disturbances.

Each of these functions can be performed in principle either in a distributed or in a centralized manner. We will call a technology “distributed” if it performs at least one function in a distributed manner. Let us consider the underlying mechanisms in more detail.

*Distributed input* means that the information is gathered in parallel by a variety of human and/or artificial agents. For example, the spread of a global epidemic can be tracked by having doctors worldwide enter new cases into a shared database the moment they observe them. Their observations may be complemented by automatic monitors that track e.g. the temperature, heart rate and blood pressure of the hospitalized patients. A simple computer program can analyze these data to show the present dispersal of the disease on a world map, and to establish typical patterns in the disease progression, so that doctors can get an idea in which stages they may have to cope with fever or a possibility of cardiac arrest.

This first processing of the input information, which might happen in a centralized system, can now be combined with a very different type of information, such as the records of phone conversations, flights, or public transport journeys. This requires not only a very different source of input, but also a very different method of processing that will typically be performed on a different system. While the first system computes the static distribution of cases, the second one computes the pattern of movement of people. Their integration in a third system then makes it possible to extrapolate where the next cases of infection are most likely to appear. The final interpretation will moreover rely on a variety of human experts who complement the results of computer algorithms with their intuitions and experience. This illustrates *distributed processing*: different human and/or artificial agents simultaneously work on (different aspects of) a common problem, while pooling their results into an overall solution.



The next stage in our example is distributed output: the prediction of the most likely future outbreaks and what to do about them now needs to be communicated to the agents involved. These will typically be medical personnel in a variety of locations. The distributed system can provide each of them with guidelines, which may be adapted to the individual (e.g. different for doctors and for nurses) and to the local circumstances (e.g. different for countryside and for city locations). These guidelines may include the most likely subjects to carry the germ (e.g. elderly people arriving from Bangkok), the symptoms to look out for, the most likely progression of the disease, and the recommended treatment for each stage of the illness, as based on the experiences up to now of various doctors worldwide. The human guidelines can be complemented by automatic programs downloaded to monitoring equipments that spot typical configurations of symptoms that tend to announce a crisis.

More generally, *distributed output* means that the system sends recommendations on how to act to a variety of human and/or artificial agents in different parts of the world, who then perform the action appropriate for their local situation. Instead of a single centralized decision, various local decisions are made and executed, albeit in a coordinated manner.

The final stage of *distributed feedback* consists of two mechanisms:

- (1) correcting for unforeseen (external) disturbances. This requires additional actions to compensate for the deviations created by the disturbances.
- (2) correcting for (internal) mistakes in selecting the right actions. This is the process of *learning* in which the information processing system becomes better by adapting its organization according to its experience.

The first mechanism can be implemented simply by feeding the new situation, as affected by the output, back to the input and processing stages. Assuming that all these stages are distributed, feedback will be distributed as well: a variety of input sensors will register the aggregated results of a variety of output effectors as affected by external events. The

second mechanism is more subtle, as it requires some degree of “rewiring” of the connections between the distributed components: some links need to become stronger, others weaker. Distribution in this case means that a variety of links are affected by the aggregate effects of a variety of actions. A “neural” mechanism for this is proposed in our mathematical model of the Global Brain (Heylighen et al., 2012): links that produce benefit are reinforced, the others weakened.

Let us illustrate such distributed feedback with the example of the global disease monitoring system. A disturbance could be an outbreak of infections in a location as yet unknown by the system. The counteraction is simply to analyze the data and send the appropriate recommendations for treatment and containment to the agents in that location. Learning in this case means updating the map of disease spread and establishing a direct link between the medical personnel in these locations and the rest of the medical community and system involved in tackling the epidemic.

### **3. Enabling Technologies**

To implement such distributed intelligence systems, we need supporting technologies. These may not be intelligent in themselves, but they provide the infrastructure necessary for an effective distribution and coordination of intelligent activity. Let us survey some basic enabling technologies that either already exist or are being developed.

#### **3.1. Distribution of information**

The world-wide distribution of information was the original function of the *Internet*: by hooking up to this “network of networks”, any computer in any place could in principle send information to any other connected computer in a quick and reliable manner. This was made possible by the TCP/IP protocol for transmitting information between computers (Leiner et al., 2009). This protocol breaks up the message into small packets that

each find a route from the sender to the receiver via a number of intermediate computers. If some packets get lost on the way, the sender is warned and immediately resends them, possibly via a different route. Thus, traveling information is distributed across different nodes and routes of the TCP/IP network.

This makes the communication particularly robust, since it will continue to function even if several of the intermediate nodes are unavailable, because of malfunction, overload or any other disturbance. The protocol was designed so that it could survive a nuclear war that would take out the major centers of communication. This decentralized, adaptive design is part of what makes the Internet such a flexible medium for information distribution, allowing programmers to build a variety of more complex systems on top of it, without having to worry about the underlying communication flows.

This infrastructure has been widely available in most developed countries since about the year 2000, but is still spreading into some of the more remote regions. This further spread is facilitated by the explosive development of *wireless* communication protocols, which make it easy to bridge many physical gaps without needing to install cables. More importantly, wireless technologies provide ubiquitous access to information, even while moving around. This enables more interactive applications, such as location-aware navigation or recommendation.

The *World-Wide Web* is a layer of protocols (HTML, HTTP, and URL) built on top of the Internet (Berners-Lee & Fischetti, 1999; Heylighen, 1994). It allows a transparent distribution of information storage: webpages on one computer contain hyperlinks to other webpages that can be situated on any computer anywhere in the world. By clicking on the link, the user is “transported” from the one computer to the other, without needing to be aware of the location of the information. Thus, complex networks of linked data can be built that are distributed over a variety of servers. This approach is developed further in *cloud storage* technology, where the distribution is so flexible and efficient that the user cannot even find out anymore exactly where the data are stored. The advantage is transparency and robustness: if a particular server is not available for whatever reason, a copy of the data

will be retrieved from another server, without the user having to worry about what is stored where.

### **3.2. Distribution of knowledge**

Information is not yet knowledge: available data first need to be processed in order to extract more general and reliable concepts and rules. This is typically done collaboratively by a variety of human experts, supported by ICT. This is the domain in which the world-wide web seems to have made the largest contributions to date, by making accepted knowledge, observations, new ideas, and plausible hypotheses freely available for discussion, analysis and reorganization.

Perhaps the best existing example of a “Global Brain”-like system emerging from such worldwide discussion is *Wikipedia*, the Internet encyclopedia that is being read and written by millions of people (Jemielniak, 2014; Kittur & Kraut, 2008). The input is clearly distributed as thousands of people are simultaneously, but independently, editing different Wikipedia articles, in order to extend, update or correct the facts that they contain. This is a very efficient way of harvesting, organizing and publicizing the collective knowledge of humanity. It can be seen as a contemporary implementation of the old ideal that H. G. Wells called the “World Brain” (Wells, 1937; Rayward, 1999). Input not only comes from people, but from a variety of software agents called “bots”, which perform various housekeeping tasks, such as correcting formatting errors, adding links and references, and gathering related articles together.

In this case, input and processing cannot really be separated as much of the input consists of reformulation or reformatting of existing material. The output too is distributed as all people with an Internet connection can (and most likely will) consult the material, while being influenced in their actions by what they find there. For example, a person suffering from a disease such as gout will find in Wikipedia a detailed exposition of the most important things known about the causes, outlook, and possible treatments for that disease, while getting suggestions about the kinds of foods or circumstances to avoid (e.g. alcohol, red meat, and

cold), but also about the ones that may help mitigate the disease (e.g. coffee, water, and heat). Output also goes to software agents, who e.g. produce statistical analyses, maps and taxonomies starting from the material in Wikipedia.

The *Semantic Web* is an emerging set of protocols to code such knowledge in an explicit format that is understandable and dependable not just for people, but for computer programs (Berners-Lee & Fischetti, 1999). These programs can make non-trivial inferences, thus allowing them to answer questions for which the answer is not written down as such, but derivable via logical inference. For example, while an article about penguins may note that penguins are birds, it is unlikely to also state that they are warm-blooded. A standard “inference engine”, on the other hand, should be able to correctly answer the question “Are penguins warm-blooded?” by combining the knowledge that penguins are birds with the knowledge that birds are warm-blooded.

### **3.3. Distribution of material objects**

Distributed intelligence not only steers the development of knowledge, but the development of artifacts. Several existing or planned technologies support this function. The success of Wikipedia is due in part to its “open access” or “open source” philosophy, according to which individual or collaborative information products can be freely used and modified by anyone (Heylighen, 2007b; Weber, 2004). The advent of *3D printers* makes it possible to extend that philosophy to physical products: a design for a material object (e.g. a bottle opener, a machine component, or a decorative vase) can be published on the Internet, downloaded and if necessary adapted by a person interested to use it, and then “printed” straight into the correct three-dimensional shape so that it is ready to use (Lipson & Kurman, 2013). Even when no 3D printer is available, many designs are made especially so that they are easy to realize with commonly available materials and components. This inspires people across the world to again start making things themselves, instead of buying industrially assembled artifacts, a trend that has been called the “*maker movement*” (Stangler & Maxwell, 2012).

The advantage is that objects can be literally made to measure, satisfying the highly specific preferences of each user, while still being much less expensive than purchased goods. Moreover, specialized parts and tools become available in places (like villages in developing countries) where market supply rarely reaches, or is prohibitively expensive. Again, input (all the designs created by people, individually or collaboratively, supported by software), processing (the many variations, improvements and extensions of available designs), and output (the many artifacts built from these designs) are all distributed. The servers that store, list and organize the available designs function here as the medium of coordination between the agents involved.

### **3.4. Distribution of energy**

Not just the production of informational and material goods, but the production of energy can be distributed in an intelligent manner. This is the idea behind various projects to build an adaptive electricity network, sometimes called *Enernet*, or *smart grid* (Li et al., 2010; Massoud Amin & Wollenberg, 2005). In the traditional approach, electricity is produced in a centralized manner, by a few large power plants running on oil, coal, or nuclear power, and then distributed via a network of electrical cables to millions of consumers. In the new approach, there are potentially millions of producers, since any company or household that owns solar panels, a windmill, watermill, or some other small-scale energy-supplying installation can add the generated electricity to the network.

Physically, this is not a problem, since the same cables can be used to both extract and inject electrical current. The problem is one of coordination: these independent producers will produce variable amounts of energy at variable times (e.g. solar panels only during the day, windmills only when there is wind), independently of the demand from the consumers (e.g. more demand in the evening because of electrical lighting). Therefore, supply will not in general match demand. This means that there is a risk of shortages or overloads at crucial moments, and a general waste of capacity whenever demand is low.

Part of the solution is in providing reservoirs able to store unused energy. But the most important part will be in implementing distributed intelligence. This may gather information from a variety of sensors (e.g. for wind speed, solar strength, temperature...) and sources (e.g. weather forecasts and statistics on energy use) in order to estimate present and future supply and demand. On that basis, the system can decide to power up or power down local energy-producing or energy-consuming installations, or adjust the price of energy according to local demand, so as to most efficiently allocate the load across the grid. The benefits of such a smart grid are obvious: less waste, less need for non-renewable, polluting resources, less risk for blackouts because a particular power plant or high-voltage line is overloaded, and more incentives for ordinary citizens to invest in clean, efficient energy.

### **3.5. Distribution of action**

Regulating the movement of energy can be extended to the movement of physical objects. This is the idea behind the *Internet of Things*, an emerging network standard for monitoring and controlling the position and state of objects such as machines, devices, goods, and building materials (Atzori, Iera, & Morabito, 2010; Welbourne et al., 2009). It suffices to equip each object with a RFID tag or other small device that can be wirelessly consulted in order to quickly get an overview of what is located where. This is particularly useful for controlling inventory, logistics and factory assembly lines, an application domain called the *industrial Internet* (Bruner, 2013), which is expected to significantly reduce the costs of production and distribution of goods. The implication is that not just people and computers but even simple material things will keep in touch via the network, thus effectively becoming part of the global brain.

Some objects will not only be able to send information about their state or location, but to receive and execute commands that tell them what to do. For example, your home thermostat may receive a message that you are on your way, and start heating so that the temperature would be pleasant by the time you arrive. Similarly, your

coffee machine may already prepare an espresso, while the blinds can open in order to let in the sunlight. The different tools in a complex, automated factory may similarly receive continuously updated commands about which items to produce in which quantities, and how best to assemble them, so as to optimally take into account the varying demands of the clients. The produced goods may leave the factory via remotely controlled transport systems, making sure the right goods reach the right clients.

Tools that affect the outside world function as “actuators” or “effectors” for the global superorganism: they convert information (commands received) into physical action. The most sophisticated effectors, which can autonomously move and manipulate objects while using sensory feedback, are commonly called “robotic”. Examples are industrial robots, self-driving cars that can be controlled remotely, autonomous submarines that explore oceanic depths, robotic arms that perform operations while being directed by a surgeon, and the drones used by the military for monitoring and attacking enemy forces. Of these technologies, drones in particular appear very promising given that they are relatively easy to build and operate, while having the great advantage that they can move in three dimensions, and thus reach vantage points inaccessible to humans (Krajník, Vonásek, Fišer, & Faigl, 2011). For example, a swarm of small flying drones can quickly survey a remote region, e.g. to locate a lost person, detect forest fires, or map the effects of flooding. Such tools make it easy to distribute sensing (input) and acting (output) capabilities across large and complex spaces.

Processing and feedback too can be distributed in this way, by letting the “robots” communicate locally with each other via wireless connections, in order to coordinate their perceptions and actions by means of self-organization (Baldassarre, Parisi, & Nolfi, 2006). An additional benefit of such self-organizing swarms of sensors and effectors is that they can produce “ad hoc” communication networks, locally propagating information from agent to agent until it reaches the fixed Internet, from where it can communicate with the rest of the world (Dressler, 2008; Elmenreich, D’Souza, Bettstetter, & de Meer, 2009). Such wireless “mesh” networks (Akyildiz, Wang, & Wang, 2005)



provide improvised, yet robust communication in complex situations, such as war zones, jungles, or crowds of people with mobile phones.

### **3.6. Distribution of identity**

Ideally, the global brain should be able to tailor its messages to each individual. For example, in the spreading epidemic scenario, different recommendations should be sent depending on whether the recipient is a doctor-specialist, a nurse, an immigration official, or a (potential) patient. In the case of a patient, the recommendations should ideally take into account the full medical record and health state of that person, so that e.g. the patient would not be suggested to take drugs that interfere with other drugs the person is already taking.

To ensure that such messages do not get to the wrong person, the distributed intelligence system should be certain about the identity of the recipient. At present, this is normally achieved by letting the user login to the system with a user name and password. This is highly inefficient, as users need to type in their data each time again in different systems, while trying to remember dozens of different passwords. Moreover, it is insecure, as people tend to use easily memorable passwords or write them down in lists, so that hackers have little difficulty in guessing or harvesting other people's passwords. This enables the novel crime of *identity theft*, in which person A pretends to be person B, thus being able to exploit person B's property or privileges (e.g. by buying goods with B's credit card details, or by posting offensive messages in B's name).

A key enabling technology for the future global brain will be a universal, secure standard for unambiguously establishing a person's identity. Several, albeit uncoordinated steps have already been made in order to create such a standard, including web-enabled electronic ID-cards in several European countries, the OpenID standard (Recordon & Reed, 2006), and ORCID (Evrard, Erdmann, Holmquist, Damon, & Dietrich, 2015), an attempt to ensure that publications are attributed to the right author. The reasons why standardization is slow to emerge tend to be social, economic and political rather than technological, as different corporations, governments and organizations are not inclined to

exchange the valuable information they hold. An additional obstacle is people's legitimate fear for invasion of privacy and abuse. However, without universal regulation, abuse of private information by hackers, corporations or governments is more rather than less likely, as no one knows who has access to which personal information, and as hardly any laws exist that specify what organizations can and cannot do with the information they possess.

Technologically, it is perfectly possible (albeit non-trivial) to develop secure schemes that anonymize data so that only the ones that really need information about an individual can get access to the specific data they require, and to nothing else. For example, a doctor who finds you collapsed in the street should be able to consult your medical record, and to send a message to your next of kin, but should not have access to your financial record. Your bank, on the other hand, should know the transactions made from your account, but not your state of health.

Next to the technological challenge, the larger challenge will be to institute a system of rules and laws that specify exactly who can use which information about a person. This system should be perfectly transparent to the individual, so that you can find out exactly what happens with your data, and have the right to withhold information that is not crucial to the functioning of an organization. The general principle is that you should be able to act anonymously for any non-crucial transaction, but that the distributed intelligence system should be able to maximally extract the collective (anonymous or non-anonymous) information that will help it to make better decisions, while being able to securely and transparently address a specific individual with personalized recommendations. Once such a computational and legal technology is in place, interactions across the Internet are likely to become much safer and more efficient.

#### **4. Satisfying needs with DIT**

Now that we have a general idea of what distributed intelligence technologies (DIT) are capable of, we need to understand what they are

most likely to be used for. Normally, an artifact is adopted because it fulfills some need, i.e. its users can achieve something valuable with the tool that would have been more difficult (or impossible) to achieve without it. For example, a refrigerator makes it easy to store food safely and thus to ensure that fresh food is available whenever needed. Note that we here use “need” in the psychological sense, as something desirable that makes us feel and function better, not in the strictest sense as something without which we would not be able to survive. For example, it is possible to live without sex, without a home, or without friends, but life would be pretty miserable in those circumstances. Therefore, whenever some innovation appears that makes it easier to satisfy a need, people will systematically (though not necessarily universally) tend to adopt that invention.

In the case of distributed technologies, there are in a sense two types of users: the local individuals, and the global system or “super-organism”. Given the size and complexity of human society, the super-organism as a whole needs to survive and thrive in order for the individual humans to survive and thrive. Without the sophisticated agricultural, industrial, economic and social infrastructure of modern civilization, our planet might at most sustain a few million human individuals, rather than the billions that it harbors now. Global problems, such as famines, wars, pollution and poverty, can be seen as remaining malfunctions or shortcomings of the superorganism. Their solution will require a better technological and institutional infrastructure and organization, something that the emerging global brain effectively facilitates (Heylighen, 2015a, 2016).

Tackling global problems in particular requires *coordination* between all the stakeholders involved in these problems. Coordination (Crowston, Rubleske, & Howison, 2006; Heylighen, 2013) mediates between humans, artifacts, and the collective systems that they form. It ensures that the different local needs do not come into conflict at higher levels, e.g. like when one person’s need for entertainment clashes with his neighbor’s need for peace and quiet. It also ensures that individual needs do not oppose global needs, as when local demand for fish undermines the sustainability of global fish stocks. Distributed intelligence can greatly facilitate coordination, via mechanisms such as

stigmergy (Heylighen, 2007b, 2008, 2015b), challenge propagation (Heylighen, 2014a) and offer networks (Goertzel, 2015; Heylighen, 2016).

As a complement to these existing studies on how DIT can satisfy global and coordination needs, the present paper will focus on the needs of individual people. For this, we can rely on classic psychological theories of motivation, such as Maslow's need hierarchy (Maslow, 1970), and on empirical research about the conditions of happiness (Heylighen & Bernheim, 2000; Veenhoven, 1997). After a brief definition of each need, we will try to imagine how present or future DIT may satisfy that need in the most efficient manner. This will provide us with a systematic list of likely future applications.

#### **4.1. Material needs**

Individuals first of all need material resources to survive. These include food, water, and oxygen, but also various raw materials, as well as the tools, clothes, shelters etc. manufactured with them. Extraction of raw materials from the environment is a function that is better addressed at the global level, since virtually no people still gather their own food, wood or wool, while sustainability requires highly coordinated action. However, individuals still need to acquire the right resources at the right time and place. In our present consumption society, this is typically done by visiting a shop, collecting the desired goods, and bringing them back home. The associated journey wastes time, space, energy, and attention, as illustrated by the ever-present curse of traffic jams.

A by now well established DIT application is Internet shopping: selecting goods from a web catalog with the help of a personalized recommender system and having them delivered wherever they are needed. This still requires travel for the delivery truck, but if distributed orders are combined in an intelligent manner, the route of the truck can be optimized for minimal cost in time and energy. The DIT application centered on 3D printing eliminates even this cost, as it removes the need for material goods to travel (Lipson & Kurman, 2013). However, even the most sophisticated future printers will still need raw materials, and

are unlikely to produce certain goods such as food. A conceivable solution is to develop an automated physical distribution network, which e.g. uses self-driving vehicles, drones or underground tubes to shuttle basic goods and materials from producer to consumer (Heylighen, 2007a).

Such DIT applications will not only change transport and production of goods, but also their economics. Generally speaking, as their production is further automated, material goods become ever less expensive. This trend will only accelerate with technologies such as 3D printing, and eventually nanotechnology (Drexler, 2013). Moreover, the “open access” model (Heylighen, 2007b) of sharing designs will completely change the price structure, making a wide range of goods practically free. Future renewable energy technologies supported by a smart grid are likely to have a similar effect on energy prices. Eventually, the resulting intelligent network may do away with money altogether, satisfying needs by directly matching offers with demands—a vision that has been called the peer-to-peer economy, sharing economy, gift economy, collaborative commons, or offer network (Goertzel, 2015; Heylighen, 2016; Rifkin, 2014). This lets us envisage an age of true abundance, where poverty or scarcity no longer exist (Diamandis & Kotler, 2012; Heylighen, 2015a).

#### **4.2. Health needs**

Our need for health, in the sense of absence of disease, is obvious enough. But health can be defined more positively as physical fitness, strength, and quality of life, i.e. an optimal state for our body characterized by a sense of well-being, long life and the capability to take on a variety of challenges (Heylighen, 2014b).

At present, health care requires an extremely expensive infrastructure that includes medical personnel, hospitals, and pharmaceutical firms. An important part of that cost is because the medical research to elucidate the causes of diseases and to develop new treatments is very slow and inefficient, requiring the collaboration of thousands of researchers, doctors and patients. A DIT approach likely to

accelerate the process is *data mining*: finding patterns in massive databases that contain billions of medical observations (Yoo et al., 2011). Eventually, all medical files maintained by doctors, hospitals and nurses will become available on the Internet, in a format suitably protecting privacy. As people start to use more and more Internet-connected sensors, these files will not just contain doctor-prescribed diagnoses and treatments, but a variety of health indicators, such as heart rate, blood pressure and temperature, as well as lifestyle elements, such as diet, exercise, exposure to sunlight, etc. There already exist a variety of smartphone apps that keep track of heart rate, running speed, temperature, altitude and other variables during their user's training sessions (Mandl, Mandel, & Kohane, 2015).

Data mining algorithms will be able to explore these "big data" collections to find correlations between various symptoms, lifestyle elements, and diseases. This will help researchers to discover new syndromes, or to establish causal connections, e.g. between a particular diet and the risk of developing a particular disease. Moreover, the distributed database will allow researchers to quickly test the effectiveness of a particular treatment by immediately determining which portion of patients actually improved with that treatment, and in what ways.

Another source of health costs is medical consultation with a highly trained expert. As expert knowledge, complemented by patterns derived from data mining, is gradually converted into computer programs, diagnosis can become largely automated. Such an expert system (e.g. The Analyst™ ([www.diagnose-me.com](http://www.diagnose-me.com)) or IBM's Watson (Ferrucci, Levas, Bagchi, Gondek, & Mueller, 2013)) would start by asking the patient a variety of initially general, then increasingly focused questions while combining the answers with publicly available data about that patient. On that basis, it would then infer and suggest the most likely diagnoses and treatments. Only in the most difficult cases would this diagnosis need to be confirmed by one or more people, thus saving many human-hours of highly qualified work. Moreover, the computer diagnostician can take into account many more factors than any individual person could, thus making it more likely that all potentially relevant tracks are covered.

Finally, DIT systems can monitor and support the treatment regime, by ensuring that the necessary medicines are taken and expected recovery stays on course. In case of emergency, the system can locate and warn the nearest doctor or ambulance, while immediately transmitting all the vital data as monitored. In more ordinary circumstances, such systems can mobilize the patient to improve his lifestyle, e.g. by stimulating the person to do the right kind of exercises and to eat the right kinds of food (Heylighen, Kostov, & Kiemen, 2013; Intille, 2004). Here too a variety of smartphone apps are being developed and quickly adopted (Kratzke & Cox, 2012). This last application will be most important for disease prevention, and most generally for maximizing overall fitness. It may well have the largest impact on medical costs, simply by preventing the most common and costly chronic diseases, such as type II diabetes, that presently plague civilization (Carrera-Bastos, Fontes-Villalba, O'Keefe, Lindeberg, & Cordain, 2011).

#### **4.3. Safety needs**

Another basic need is feeling safe from danger. Apart from the health problems discussed before, the most common dangers in our present society are accidents and crime. Accidents are most frequently caused by traffic, followed by working with machinery. The spread of sensors and self-regulation in vehicles, buildings and machines will make it increasingly easy for DIT applications to prevent potential accidents. For example, sensors that keep track of the positions and speeds of the cars in a lane can send a braking signal to the motor as soon as the preceding car slows down too much, thus avoiding a collision. Google has built driverless cars that combine sensors such as these with a wireless connection that allows the vehicle to plan its journey taking into account road maps and traffic conditions (Hars, 2015; Luettel, Himmelsbach, & Wuensche, 2012). These self-driving vehicles are reputedly already safer than human-driven ones.

Sophisticated sensors and control programs can similarly reduce the risk of crime, e.g. by using biometric data such as iris scans, fingerprints, and voice recognition to identify an individual, and thus

make sure that only the rightful owner of a car, smartphone or house gets access to it. If access cannot be fully secured, sensors aided by smart algorithms can warn the police whenever a suspicious activity (such as the breaking of a window or an intrusion) takes place. Distributed intelligence relies on more than such devices, though: it makes full use of human eyes, ears and brains. Smartphone technology already makes it easy for people to immediately warn the police whenever they witness a crime, or any apparent preparation for it, including detailed data such as place, time, and video recordings. Distributed databases, such as Ushahidi, moreover make it easy to share, organize and visualize such data, so as to produce an overview of what is happening where. This makes it easier to plan and coordinate interventions in confuse situations such as fires, riots, or natural disasters (Gao, Barbier, & Goolsby, 2011). Self-organizing sensor networks may even provide ad hoc guidance in situations where all traditional communications have broken down, such as an explosion in a tunnel (Dressler, 2008). After the events, data mining on the collected information should allow uncovering the most common causes of accidents, violence and crime, and thus help to formulate guidelines and precautions for preventing them.

#### **4.4. Social needs**

Being able to rely on others is another fundamental human need. To be truly happy, everybody needs friends, lovers or family, and the feeling of belonging to a community. Such relationships are easily supported by DITs.

At present, such social applications are perhaps the most popular ones of all, as illustrated by the explosive development of “*social media*” or “social networks”, such as Facebook, Twitter and LinkedIn, together with a seemingly endless list of Internet communities, forums and discussion groups. Critics have a point in observing that social interactions sustained in this way are often superficial compared to the more traditional ones. But this is most likely due to the novelty of the medium, and the fact that people like to play and explore the many new functions, without as yet having learned to distinguish services that



satisfy true needs from gadgets that produce cheap thrills. On the other hand, since its early days the Internet has been used to create and sustain deep personal and professional relationships across distances too large to allow for face-to-face contacts (Baker, 2005).

The most obvious DIT application is to use the variety of data available about people worldwide to suggest good “matches”, i.e. others that you might like to get to know, as professional partners, friends, or potential lovers. Relevant data include existing social connections, interests, personality traits, age, location, and professional, educational, social and cultural background. Matchmaking sites are one of the great success stories of the web (Whitty, Baker, & Inman, 2007). According to surveys, in one third of recent American marriages the couple met online. Such relationships started online even appear slightly more successful than the traditional ones (Toma, 2015). With more data becoming available and smarter recommendation algorithms, the effectiveness of online dating can only continue to increase.

Moreover, with the spread of high-bandwidth video communication and future, more sophisticated sensing methods (e.g. virtual reality, emotion recognition, telepresence), remote interaction becomes increasingly more similar to face-to-face interaction. This will help people to communicate subtle feelings across the net, and thus experience a true emotional connection with their communication partners.

Finally, the Internet has an impressive track record in establishing new *communities*, i.e. groups of people with a shared interest exchanging information and doing things together. People with unusual interests (e.g. those suffering from a rare disease, having an uncommon hobby, or non-standard sexual preferences) in the past would have felt lonely and alienated in their local environment. Now these people can easily find like-minded people who will not only accept them for what they are, but provide them with encouragement, feedback and support. The effect on personal well-being can be profound, as people no longer feel excluded, but rather appreciated for their contributions to the community, while always having someone available to give advice or help.

#### 4.5. *Achievement needs*

To be truly happy, people need to develop a sense of achievement or mastery, of feeling that their actions are effective and that they have some degree of power or control over their situation. Moreover, people like to be recognized by others for such achievements, so that they can enjoy a sense of esteem, respect or status.

One of the reasons for the popularity of Internet communities and social media is precisely that they make it easy for people to get recognition for their contributions from peers, thus allowing them to build up a good reputation. The popularity of games, on the other hand, is due for an important part to the fact that they provide immediate feedback about the actions performed by the player, in the form of points, scores, and awards (Heylighen et al., 2013). This creates a sense of achievement, as players can graduate into increasingly advanced “levels” of expertise in the game. While these gaming applications may seem rather frivolous, the underlying psychological mechanism of *flow* makes a real contribution to people’s level of happiness (Chen, 2007; Cowley, Charles, Black, & Hickey, 2008; Nakamura & Csikszentmihalyi, 2002). Moreover, this mechanism can be extended by techniques such as “gamification” to tasks that are very serious indeed, such as studying mathematics, or increasing fitness (Deterding, Sicart, Nacke, O’Hara, & Dixon, 2011).

The concept of *mobilization system* (Heylighen et al., 2013) refers to a new type of DIT that encourages and helps people to work towards a worthwhile objective, thus boosting their sense of achievement. In addition, a number of Internet community forums (e.g. Stack Overflow; Mamykina, Manoim, Mittal, Hripcsak, & Hartmann, 2011) are developing tools for measuring individuals’ overall contribution, thus automating the development of reputation (De Alfaro, Kulshreshtha, Pye, & Adler, 2011), and rewarding helpful people for their good work.

#### 4.6. Cognitive needs

People have an innate desire to know, to learn and to understand the world around them. Providing such knowledge is another one of the best-established uses of DIT. Search engines, such as the one of Google, have learned to anticipate and understand the most common questions that people have, and to find the documents most likely to provide a high-quality answer. Wikipedia harnesses the collective intelligence of all people to produce standard accounts of all knowledge domains about which some degree of consensus exists (Jemielniak, 2014), while the Semantic Web makes such knowledge available to machines as well as to people.

Both Wikipedia and the Semantic Web will continue to grow in the size and the quality of the knowledge they cover, until they contain all the knowledge discovered by humanity. At that moment, any question someone may ask for which an answer is possible given the present state of knowledge will receive that answer immediately. This is the practical equivalent of *omniscience*, a property that was hitherto considered an attribute of God (Heylighen, 2015a; Otlet, 1935).

The most commonly used knowledge should not just be available via the network, but as much as possible inside people's own brains, so that they immediately can apply it to the situation at hand. This can be achieved by *learning*. DIT are starting to revolutionize the process of education, thus making learning much easier, more efficient and more enjoyable (Heylighen et al., 2013). One of the advantages is that learning can be perfectly tailored to the individual, by providing the right kind of challenges at the right moment, depending on the interests and capabilities of the learner, and not on the time or place where that learner resides. This is already achieved to some degree by "Massively Online Open Courses" (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).

A straightforward extrapolation of these developments tells us that in the near future, people anywhere in the world will be able to study and obtain degrees on any subject, from basic literacy and numeracy to PhD level science, by freely following inspiring course materials

remotely, while getting feedback from automatic evaluation programs, peers studying the same material, and—if need be—teachers. This will boost 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 at least a Masters degree in an advanced domain, while having a broad and deep general education covering a variety of common and uncommon topics...

#### **4.7. *Self-actualization or growth needs***

At the highest level of Maslow's hierarchy of needs is the desire to maximally develop one's own capabilities, i.e. to realize one's potential, or to grow psychologically (Heylighen, 1992; Maslow, 1970). This need is the final one in the sense that it has no endpoint, as one can always grow wiser or more capable. It ensures that no one will ever get bored because of having achieved all her desires.

The Internet can be seen as a space for unlimited exploration, allowing individuals to discover new opportunities, to take on new challenges, to develop new insights, to learn new skills, and to meet new people (Last, 2016). Indeed, new resources join the Internet at a much faster pace than any individual could explore them, thus ensuring that there is always much more to discover than what people already know. In that sense, the Internet allows an unlimited growth in knowledge, social connections, resources and wisdom, even for people who are too old or too frail to leave their homes.

Moreover, distributed intelligence algorithms will recommend activities or domains to explore that are tailored to an individual's personal interests and competences, while challenging them to go ever further in their endeavors. The path of learning can even be optimized by taking into account the experiences of others that followed a similar path before (Gutiérrez, Valigiani, Jamont, Collet, & Delgado Kloos, 2007). This ensures that people are constantly stimulated to develop themselves further and to take on more difficult challenges, so that personal development never needs to come to an end.

## 5. Conclusion

This paper has tried to survey the major present and future applications of what we have called “distributed intelligence technologies”. These are systems that tackle complex challenges in a distributed manner, by collecting input from many different sources, storing and processing that information simultaneously in different places, and sending their outputs to different human or technological agents working in parallel.

To achieve this, such systems need a global communication network, such as the Internet, to interconnect all people, computers and machines that are involved in solving the problem. Moreover, they need to make sure that the inputs and actions of all these agents are efficiently coordinated, so that they produce a coherent result. Supporting technologies for such distributed problem solving, such as the Semantic Web, the Internet of Things and the Smart Grid, are being developed at a rapid pace. They at least need to enable the efficient distribution of information, knowledge, physical objects, energy, actions and personal identity. The integration of all these technologies is expected to produce a distributed intelligence at the planetary scale, the *Global Brain*.

The capabilities of a Global Brain in tackling the problems of our planetary society have been surveyed elsewhere (Heylighen, 2007a, 2015a, 2016). The present paper has therefore focused on its applications at the level of individuals: how can distributed intelligence technologies satisfy our basic human needs? By extrapolating from existing applications, such as social media, data mining, online shopping and computer-supported education, we have shown that these technologies can in principle fulfill our material, health, safety, social, achievement, cognitive, and even growth needs. Once we have learned to deal more effectively with these technologies—and in particular some of their stress-producing side effects, such as information overload, distraction, and unpredictability (Heylighen, 2015a; Heylighen et al., 2013)—this should greatly facilitate personal fulfillment and well-being.

## References

- Akyildiz, I. F., Wang, X., & Wang, W. (2005). Wireless mesh networks: a survey. *Computer Networks*, 47(4), 445–487. <http://doi.org/10.1016/j.comnet.2004.12.001>
- Ashby, W. R. (1956). Design for an intelligence-amplifier. *Automata Studies*, 215–234.
- Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. *Computer Networks*, 54(15), 2787–2805.
- Baker, A. J. (2005). *Double click: romance and commitment among couples online*. Hampton Press.
- Baldassarre, G., Parisi, D., & Nolfi, S. (2006). Distributed coordination of simulated robots based on self-organization. *Artificial Life*, 12(3), 289–311.
- Ball, P. (2012). *Why Society is a Complex Matter: Meeting Twenty-first Century Challenges with a New Kind of Science* (2012th ed.). New York: Springer.
- Bechtel, W., & Abrahamsen, A. (1991). *Connectionism and the mind: An introduction to parallel processing in networks*. Basil Blackwell.
- 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>
- Bernstein, A., Klein, M., & Malone, T. W. (2012). Programming the Global Brain. *Communications of the ACM*, 55(5), 1.
- Bruner, J. (2013). *Industrial Internet*. O'Reilly Media, Inc.
- Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd, J., Theraula, G., & Bonabeau, E. (2003). *Self-organization in biological systems*. Princeton Univ Pr.
- 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. <http://doi.org/10.2147/RRCC.S16919>
- Chen, J. (2007). Flow in games (and everything else). *Communications of the ACM*, 50(4), 31–34.
- Clark, A. (1998). Embodied, situated, and distributed cognition. *A Companion to Cognitive Science*, 506–517.
- Cowley, B., Charles, D., Black, M., & Hickey, R. (2008). Toward an understanding of flow in video games. *Computers in Entertainment (CIE)*, 6(2), 20.
- Crowston, K., Rubleske, J., & Howison, J. (2006). Coordination theory. *Human-Computer Interaction and Management Information Systems: Foundations*, 120.
- De Alfaro, L., Kulshreshtha, A., Pye, I., & Adler, B. T. (2011). Reputation systems for open collaboration. *Communications of the ACM*, 54(8), 81–87.
- 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>
- Deterding, S., Sicart, M., Nacke, L., O'Hara, K., & Dixon, D. (2011). Gamification. using game-design elements in non-gaming contexts. In *Proceedings of the 2011 annual*

- conference extended abstracts on Human factors in computing systems (pp. 2425–2428).
- Diamandis, P. H., & Kotler, S. (2012). *Abundance: The Future Is Better Than You Think*. Free Press.
- 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.
- 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.
- Evrard, A. E., Erdmann, C., Holmquist, J., Damon, J., & Dietrich, D. (2015). Persistent, Global Identity for Scientists via ORCID. *arXiv:1502.06274 [astro-Ph, Physics:physics]*. Retrieved from <http://arxiv.org/abs/1502.06274>
- Ferrucci, D., Levas, A., Bagchi, S., Gondek, D., & Mueller, E. T. (2013). Watson: Beyond Jeopardy! *Artificial Intelligence*, 199–200, 93–105. <http://doi.org/10.1016/j.artint.2012.06.009>
- 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. <http://doi.org/10.1145/1133265.1133268>
- Gao, H., Barbier, G., & Goolsby, R. (2011). Harnessing the crowdsourcing power of social media for disaster relief. *Intelligent Systems, IEEE*, 26(3), 10–14.
- Goertzel, B. (2002). *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers.
- Goertzel, B. (2015). Beyond Money: Offer networks, a potential infrastructure for a post-money economy. In B. Goertzel & T. Goertzel (Eds.), *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity* (pp. 522–549). Humanity+ Press.
- Gutiérrez, S., Valigiani, G., Jamont, Y., Collet, P., & Delgado Kloos, C. (2007). A swarm approach for automatic auditing of pedagogical planning. In *Advanced Learning Technologies, 2007. ICALT 2007. Seventh IEEE International Conference on* (pp. 136–138). Retrieved from [http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=4280973](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4280973)
- Hars, A. (2015). Self-Driving Cars: The Digital Transformation of Mobility. In C. Linnhoff-Popien, M. Zaddach, & A. Grahl (Eds.), *Marktplätze im Umbruch* (pp. 539–549). Springer Berlin Heidelberg. Retrieved from [http://link.springer.com/chapter/10.1007/978-3-662-43782-7\\_57](http://link.springer.com/chapter/10.1007/978-3-662-43782-7_57)
- Heylighen, F. (1992). A cognitive-systemic reconstruction of Maslow's theory of self-actualization. *Behavioral Science*, 37(1), 39–58. <http://doi.org/10.1002/bs.3830370105>

- Heylighen, F. (1994). World-Wide Web: a distributed hypermedia paradigm for global networking. In *SHARE EUROPE SPRING MEETING* (pp. 355–368).
- Heylighen, F. (1999). Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map. *Computational & Mathematical Organization Theory*, 5(3), 253–280. <http://doi.org/10.1023/A:1009690407292>
- Heylighen, F. (2007a). The Global Superorganism: an evolutionary-cybernetic model of the emerging network society. *Social Evolution & History*, 6(1), 58–119.
- 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 (Eds.), *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. In *Globalization as evolutionary process: modeling global change* (p. 284). Routledge. Retrieved from <http://pcp.vub.ac.be/papers/AcceleratingEvolution.pdf>
- Heylighen, F. (2011). Conceptions of a Global Brain: an historical review. In *Evolution: Cosmic, Biological, and Social*, eds. Grinin, L. E., Carneiro, R. L., Korotayev A. V., Spier F. (pp. 274–289). Uchitel Publishing. Retrieved from <http://pcp.vub.ac.be/papers/GB-conceptions-Rodrigue.pdf>
- Heylighen, F. (2013). Self-organization in Communicating Groups: the emergence of coordination, shared references and collective intelligence. In À. Massip-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. (2014a). Challenge Propagation: Towards a theory of distributed intelligence and the global brain. *Spanda Journal*, V(2), 51–63.
- Heylighen, F. (2014b). Cybernetic Principles of Aging and Rejuvenation: the buffering-challenging strategy for life extension. *Current Aging Science*, 7(1), 60–75. <http://doi.org/10.2174/1874609807666140521095925>
- Heylighen, F. (2015a). Return to Eden? Promises and Perils on the Road to a Global Superintelligence. In B. Goertzel & T. Goertzel (Eds.), *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*. Humanity+ Press. Retrieved from <http://pespmc1.vub.ac.be/Papers/BrinkofSingularity.pdf>
- Heylighen, F. (2015b). Stigmergy as a Universal Coordination Mechanism I: definition and components. *Cognitive Systems Research*, (submitted).
- Heylighen, F. (2016). Towards an Intelligent Network for Matching Offer and Demand: from the sharing economy to the Global Brain. *Technological Forecasting & Social Change*, in press.
- 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. <http://doi.org/10.1023/A:1010099928894>



- 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., Heath, M., & Van Overwalle, F. J. (2004). The Emergence of Distributed Cognition: a conceptual framework. In *Proceedings of collective intentionality IV*.
- 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>
- Holland, J. H. (1992). Complex adaptive systems. *Daedalus*, 17–30.
- Hutchins, E. (2000). Distributed cognition. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences*. Elsevier Science.
- Hutchins, E., & Lintern, G. (1995). *Cognition in the Wild* (Vol. 262082314). MIT press Cambridge, MA. Retrieved from [http://books.google.be/books?id=AfupQgAACAAJ&dq=Cognition+in+the+Wild&hl=en&sa=X&ei=YjLWT6PGLaOn0QWwyqiBBA&redir\\_esc=y](http://books.google.be/books?id=AfupQgAACAAJ&dq=Cognition+in+the+Wild&hl=en&sa=X&ei=YjLWT6PGLaOn0QWwyqiBBA&redir_esc=y)
- Intille, S. S. (2004). A new research challenge: persuasive technology to motivate healthy aging. *Information Technology in Biomedicine, IEEE Transactions on*, 8(3), 235–237.
- Jemielniak, D. (2014). *Common Knowledge?: An Ethnography of Wikipedia*. Stanford University Press.
- Kittur, A., & Kraut, R. E. (2008). Harnessing the wisdom of crowds in wikipedia: quality through coordination. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work* (pp. 37–46). New York, NY, USA: ACM. <http://doi.org/10.1145/1460563.1460572>
- Krajník, T., Vonásek, V., Fišer, D., & Faigl, J. (2011). AR-Drone as a Platform for Robotic Research and Education. In D. Obdržálek & A. Gottscheber (Eds.), *Research and Education in Robotics - EUROBOT 2011* (pp. 172–186). Springer Berlin Heidelberg. Retrieved from [http://link.springer.com/chapter/10.1007/978-3-642-21975-7\\_16](http://link.springer.com/chapter/10.1007/978-3-642-21975-7_16)
- Kratzke, C., & Cox, C. (2012). Smartphone Technology and Apps: Rapidly Changing Health Promotion. *International Electronic Journal of Health Education*, 15, 72.
- Last, C. (2016). Self-actualization in the Commons. *Technological Forecasting and Social Change*.
- Leiner, B. M., Cerf, V. G., Clark, D. D., Kahn, R. E., Kleinrock, L., Lynch, D. C., ... Wolff, S. (2009). A Brief History of the Internet. *SIGCOMM Comput. Commun. Rev.*, 39(5), 22–31. <http://doi.org/10.1145/1629607.1629613>

- Li, F., Qiao, W., Sun, H., Wan, H., Wang, J., Xia, Y., ... Zhang, P. (2010). Smart Transmission Grid: Vision and Framework. *IEEE Transactions on Smart Grid*, *1*(2), 168–177. <http://doi.org/10.1109/TSG.2010.2053726>
- Lipson, H., & Kurman, M. (2013). *Fabricated: The New World of 3D Printing*. John Wiley & Sons.
- Luettel, T., Himmelsbach, M., & Wuensche, H. J. (2012). Autonomous Ground Vehicles—Concepts and a Path to the Future. *Proceedings of the IEEE*, *100*(13), 1831–1839.
- Malone, T. W., Laubacher, R., & Dellarocas, C. (2010). The collective intelligence genome. *IEEE Engineering Management Review*, *38*(3), 38.
- Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G., & Hartmann, B. (2011). Design lessons from the fastest q&a site in the west. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (pp. 2857–2866). New York, NY, USA: ACM. <http://doi.org/10.1145/1978942.1979366>
- Mandl, K. D., Mandel, J. C., & Kohane, I. S. (2015). Driving Innovation in Health Systems through an Apps-Based Information Economy. *Cell Systems*, *1*(1), 8–13. <http://doi.org/10.1016/j.cels.2015.05.001>
- Maslow, A. H. (1970). *Motivation and personality* (2nd ed.). New York: Harper & Row.
- Massoud Amin, S., & Wollenberg, B. F. (2005). Toward a smart grid: power delivery for the 21st century. *Power and Energy Magazine, IEEE*, *3*(5), 34–41.
- 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.
- Michael, D. R., & Chen, S. L. (2005). *Serious games: Games that educate, train, and inform*. Muska & Lipman/Premier-Trade.
- Miller, J. G. (1965). Living systems: Basic concepts. *Behavioral Science*, *10*(3), 193–237.
- 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, N.J: Princeton University Press.
- 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.
- Otlet, P. (1935). *Monde: Essai d'Universalisme*. Brussels: Mundaneum.
- Rayward, W. B. (1999). H. G. Wells' s idea of a World Brain: A critical reassessment. *Journal of the American Society for Information Science*, *50*(7), 557–573.
- Recordon, D., & Reed, D. (2006). OpenID 2.0: A Platform for User-centric Identity Management. In *Proceedings of the Second ACM Workshop on Digital Identity Management* (pp. 11–16). New York, NY, USA: ACM. <http://doi.org/10.1145/1179529.1179532>

- Rifkin, J. (2014). *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. Palgrave Macmillan. Retrieved from <http://www.bookdepository.com/Zero-Marginal-Cost-Society-Jeremy-Rifkin/9781137278463>
- 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>
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. San Diego: University of California Press.
- Russell, P. (1995). *The global brain awakens: Our next evolutionary leap*. Global Brain, Inc.
- Stangler, D., & Maxwell, K. (2012). DIY Producer Society. *Innovations: Technology, Governance, Globalization*, 7(3), 3–10. [http://doi.org/10.1162/INOV\\_a\\_00134](http://doi.org/10.1162/INOV_a_00134)
- Stock, G. (1993). *Metaman: the merging of humans and machines into a global superorganism*. Simon & Schuster.
- Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor.
- Thompson, C. (2011). How Khan Academy is changing the rules of education. *Wired Magazine*, (August), 126.
- Toma, C. L. (2015). Online Dating. In *The International Encyclopedia of Interpersonal Communication*. John Wiley & Sons, Inc. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1002/9781118540190.wbeic118/abstract>
- Veenhoven, R. (1997). Advances in Understanding Happiness. *Revue Québécoise de Psychologie*, 18(2), 29–74.
- Weber, S. (2004). *The success of open source*. Cambridge University Press.
- Welbourne, E., Battle, L., Cole, G., Gould, K., Rector, K., Raymer, S., ... Borriello, G. (2009). Building the internet of things using RFID: the RFID ecosystem experience. *Internet Computing, IEEE*, 13(3), 48–55.
- Wells, H. (1937). *World Brain*. Ayer Co Pub. Retrieved from <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/0836920333>
- Whitty, M. T., Baker, A. J., & Inman, J. A. (2007). *Online matchmaking*. Palgrave Macmillan. Retrieved from <http://www.palgrave.com/rights/pdfs/Cultural%20and%20Media%20Studies.pdf>
- Yoo, I., Alafaireet, P., Marinov, M., Pena-Hernandez, K., Gopidi, R., Chang, J.-F., & Hua, L. (2011). Data Mining in Healthcare and Biomedicine: A Survey of the Literature. *Journal of Medical Systems*, 36(4), 2431–2448. <http://doi.org/10.1007/s10916-011-9710-5>