

Challenge Propagation: Towards a theory of distributed intelligence and the global brain

Francis Heylighen

Global Brain Institute

Vrije Universiteit Brussel

Abstract: We sketch a foundation for a new theory of distributed intelligence, based on the process of *challenge propagation*, which extends the mechanism of spreading activation in neural networks to the collective intelligence emerging from a network of interacting agents. Challenge propagation is a form of self-organizing, distributed processing that allows agents to collectively tackle challenges too complex for a single agent, and that can be mathematically and computationally modelled. The basic idea is to combine the notion of “challenge”, which is defined as a phenomenon that elicits action from an agent, with the notion of “propagation”, which denotes the process by which such phenomenon is iteratively transmitted from agent to agent. A challenge is a generalization of the notions of problem, opportunity and activation. It can be characterized by valence (positive or negative), prospect, mystery and difficulty. An agent’s action on a challenge will typically “relax” the challenge, but not resolve it altogether, so that some degree of challenge remains for further agents to act upon. Propagation occurs either via a shared medium in which challenging traces are left for others (stigmergy), or via a network of agent-to-agent links learned through reinforcement of successful transmissions.

Introduction

Contemporary science sees societies, organisms and brains as *complex adaptive systems* (Ball, 2012; Holland, 1992; 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. 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.

The present paper will focus on the *distributed intelligence* (Fischer, 2006) of such a self-organizing system, because this is what most fundamentally distinguishes the new paradigm from the older paradigm, which sees problem solving and decision making as centralized,

sequential processes. We will define *intelligence* as the ability to process information so as to efficiently solve problems and exploit opportunities. What are considered problems, opportunities—or more generally *challenges*—will depend on the goals and values of the decision-maker, who can be an individual, an organization, or a superorganism (Heylighen, 1999, 2012b). Efficiently dealing with a challenge means selecting and performing the right actions that solve the problem or exploit the opportunity.

Traditional models of intelligence in cognitive science and artificial intelligence see the problem solving as a process of 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 divided into 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.

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—e.g. 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.

This paper wishes to introduce a new paradigm for modelling this process, *challenge propagation*, which synthesizes my older work on spreading activation in individual and collective intelligence (Heylighen, 1999; Heylighen & Bollen, 2002), and my more recent ontology of action (Heylighen, 2011b, 2013). The basic idea is to combine the notion of “challenge”, which is defined in the action ontology as a phenomenon that elicits action from an agent, with the notion of “propagation” or “spreading”, which comes from models of neural networks, memetics (Gabora, 1993; Heylighen & Chielens, 2008), and complex systems, and which denotes the process by which some phenomenon is iteratively transmitted from a point in a node in a network to the neighboring nodes.

The intention of this work is to provide a conceptual and mathematical foundation for a new theory of the *Global Brain* (Goertzel, 2002; Heylighen, 2008, 2011a, 2011c), which is defined as the distributed intelligence emerging from all people and machines as connected by the

Internet. However, the notion of challenge propagation seems simple and general enough to also provide a foundation for a theory of distributed intelligence in general. This includes human intelligence—which as neural network researchers have shown is distributed over the billions of neurons in the brain (Bechtel & Abrahamsen, 1991; McLeod et al., 1998)—, the collective intelligence of insects, but also various as yet poorly understood forms of intelligence in e.g. bacteria (Ben-Jacob, Becker, Shapira, & Levine, 2004) or plants (Trewavas, 2003).

In fact, I assume that—in contrast to traditional, sequential models of artificial intelligence—all forms of “natural” intelligence are distributed. This means that they emerge from the interactions between a collective of autonomous components or agents that are working in parallel. This perspective has also been called the “society of mind” (Minsky, 1988): a mind or intelligence can be seen as a collaboration between relatively independent agents. More generally, intelligence can be viewed as the capability for coordinated, organized activity. Excluding “intelligent design” accounts—which presuppose the very intelligence they purport to explain—this means that intelligence must ultimately be the result of self-organization (Heylighen, 2013), a process which typically occurs in a distributed manner.

Another reason to focus on distributed intelligence is that traditional intelligence models—in which a well-defined agent solves a well-defined problem (and then stops)—are completely unrealistic for describing complex adaptive systems, such as an organization, the Internet, or the brain. In such systems, everything is “smeared out” across space, time and agents: it is never fully clear who is addressing which problem where or when. Many components contribute simultaneously to many “problem-solving” processes, and problems are rarely completely solved: they rather morph into something different. That is why the notion of “problem” will need to be replaced by the broader notion of “challenge” and the sequential, localized process of “search” (for a problem solution) by the parallel, distributed process of “propagation”.

The difficulty, of course, is to represent such a complex, ill-defined process in a precise, mathematical or computational manner. There exist already a number of useful paradigms for doing this, including multi-agent systems, complex dynamic systems, neural networks, and stigmergy (Heylighen, 2015; Parunak, 2006). The challenge propagation paradigm is intended to synthesize the best features of these different models. The present paper will sketch the conceptual foundations that are necessary to build such a model, while leaving the mathematical development for another paper (Heylighen, Busseniers, Veitas, Vidal, & Weinbaum, 2012).

A brief review of intelligence models

The most simple and common definition of intelligence is *the ability to solve problems* (Heylighen, 1999). A problem can be defined as a difference between the present situation (the initial state), and an ideal or desired situation (the goal state or solution). Problem solving then means finding a path through the “problem space” that leads from the initial state (say, x) to the goal (say, y) (Heylighen, 1988; Newell & Simon, 1972). This requires determining the right sequence of steps that leads from x to y (see Fig. 1).

For non-trivial problems, the number of potential paths that need to be explored increases exponentially with the number of steps, so that it quickly becomes astronomical. For example, if at each stage you have the choice between 10 possible steps, there will be 10^n possible paths

of length n . That makes one trillion for a path of merely 12 steps long! That is why “brute force” approaches (trying out all possible paths in order to find the right one) in general do not work, and need to be complemented by what we conventionally call “intelligence”.

The more problems an agent can solve, the more intelligent it is. Note that this definition does not provide an absolute measure of intelligence, as the number of problems that a non-trivial agent can solve is typically infinite. Therefore, counting the number of solvable problems does not produce the equivalent of an IQ. On the other hand, the present definition does produce a *partial ordering*: an agent A is more intelligent than another agent B, if A can solve all problems that B can solve, and some more. In general, though, A and B are incomparable, as B may be able to tackle some problems that A cannot deal with, and vice versa.

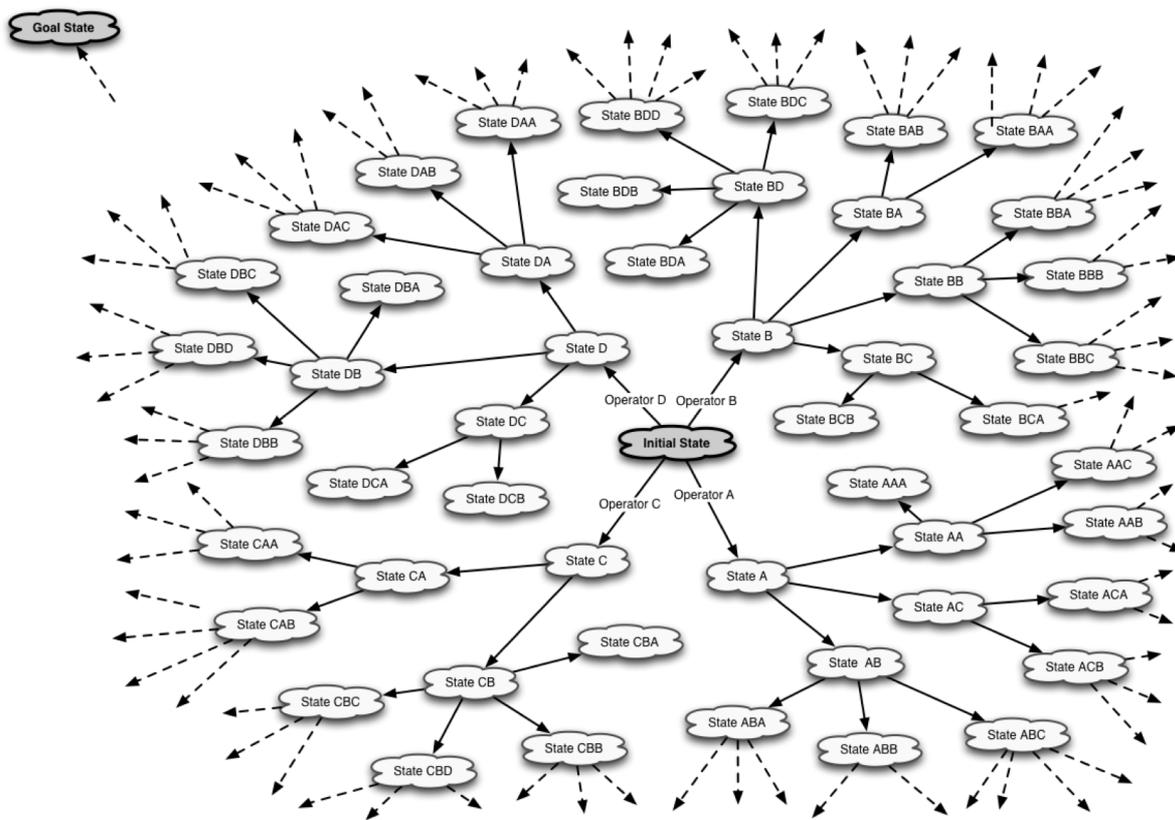


Figure 1: an illustration of the exponential explosion in the number of possible paths leading from an initial problem state via subsequent steps (or “operators”) to the goal state or problem solution.

The partial order provides us with an unambiguous criterion of progress: if an agent, by learning, evolution, or design, manages to solve additional problems relative to the ones it could deal with before, it has become objectively more intelligent. Natural selection entails that more intelligent agents will sooner or later displace less intelligent agents, as the latter will at some stage be confronted with problems that they cannot solve, but that the more intelligent ones can solve. Thus, the more intelligent ones have a competitive advantage over

the less intelligent ones. Therefore, we may assume that evolutionary, social, or technological progress will in general increase intelligence in an irreversible way.

Yet, we should remember that in practice intelligence is highly context-dependent: more important than the absolute number of problems you can solve, is whether you can solve the problems that are significant for you in your present situation. Adding the capability to solve some purely theoretical problems that have no value in your present or future environment will in general not increase your fitness (i.e. probability of long-term survival)—and may even decrease it if it would make you waste time on contemplating irrelevant issues.

The simplest model of intelligence is a *look-up table* or *mapping*. This is a list of *condition-action rules*, of the form: if your problem is x , then the (action you need to perform to attain the) solution is y . In short: if x , then y , or, even shorter: $x \rightarrow y$. An example is the table of multiplication, which lists rules such as: if your problem is 7×7 , then the solution is 49.

The next, more complex model of intelligence is a deterministic *algorithm*. This is a sequence of actions that need to be performed on the initial state in a particular order. The sequence is typically iterated until the state it produces satisfies the condition for being a solution. An example is a procedure to calculate 734×2843 or a program that determines the first 100 prime numbers. Such deterministic procedures to manipulate numbers or, more generally, lists of symbols, have given rise to the notion of intelligence as *computation*.

A deterministic algorithm (like finding prime numbers) is guaranteed to produce an acceptable solution after a finite number of steps. Problems that are more complex do not offer such a guarantee: trial-and-error will be needed, and, by definition, you do not know whether any trial will produce a solution or an error. In this case, the best you can hope for is a *heuristic*: a procedure that suggests plausible paths towards a solution. Heuristics do not necessarily produce the correct solution: they merely reduce the amount of search you would have to perform with respect to a “brute force”, exhaustive exploration of the problem space. The better the heuristic, the larger the reduction in search and the higher the probability that you would find the solution after a small number of steps.

The view of problem solving as computation or as heuristic search seems to imply a *sequential* process, in which the different actions are performed one by one in a central location. A first step in our intended generalization towards distributed processes is the reinterpretation of problem solving as *information processing*. The initial state or problem statement can be interpreted as a piece of information received by the agent. The solution of the problem is a new piece of information produced by the agent in response to the problem statement. The task of the intelligent agent is then to transform or process the input information (problem, initial state, “question”) via a number of intermediate stages into the output information (solution, goal state, “answer”).

While the term “information processing” is widespread, its meaning remains surprisingly vague: how exactly is a given piece of information transformed into a new—and presumably more useful or meaningful—piece of information? Apart from deterministic computation, which is merely a very specific case of processing, I do not know of any general, formal model of information processing. But this vagueness is an advantage as it allows us to consider a variety of mechanisms and models beyond sequential algorithms or search.

One of the most successful alternative models of information processing can be found in neural networks (McLeod et al., 1998). In the simplest case, the network consists of connected units or nodes arranged in subsequent “layers”, with the connections pointing from the “input layer”, via one or more “hidden” layers, to the final “output layer” (see Fig. 2). Information processing happens simply by presenting the information to the input layer (in

the form of a pattern of activation distributed across the nodes), letting that information propagate through the hidden layers (during which the activation pattern changes depending on the strengths of the connections), and collecting the processed information at the output layer by reading the activation pattern of the final nodes. This seems to be in essence how the brain processes information: the input layer represents the neurons activated by sensory organs (perception), the output layer represents the neurons that activate motor organs (action), and the hidden layers represent the intervening brain tissue processing the sensory information.

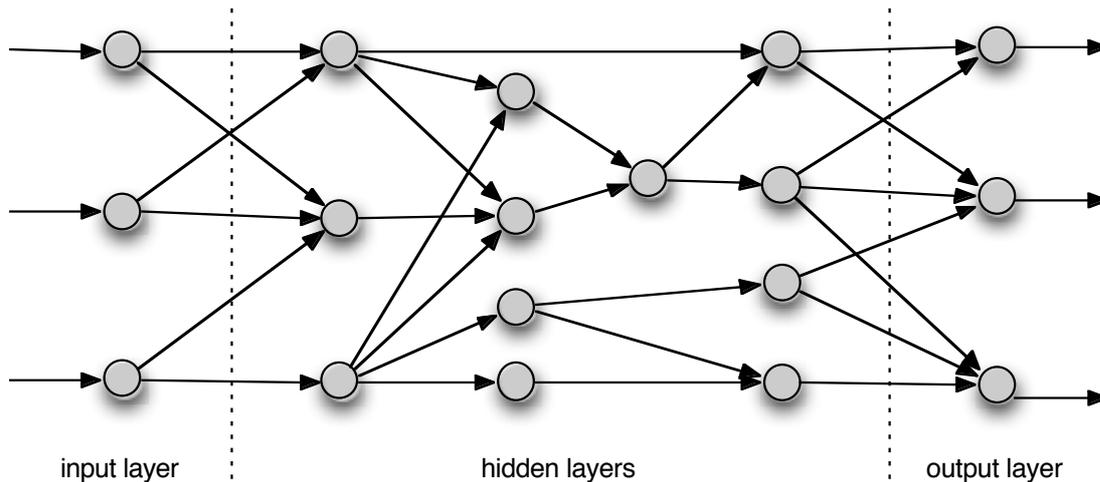


Figure 2: a neural network with links (represented by arrows) connecting nodes (represented by circles). The problem is posed by differentially activating the nodes in the input layer. This activation propagates across the hidden layers while undergoing processing. The final activation pattern of the output layer is read off as the solution.

The more general version of such a “feedforward” network is called a “recurrent” network. The difference is that a recurrent network allows activation to cycle back to nodes activated earlier. Thus, there is no imposed direction “forward”, from input layer to output layer. The input in this case is simply the initial pattern of activation over all nodes. The output is the final pattern of activation after it has settled into a stable configuration.

Compared to the sequential models of intelligence, neural networks have two big advantages:

- processing happens in a parallel, distributed manner, making it more robust and flexible;
- the network does not need an explicit program telling it how to perform the process: it can learn from experience.

The distributed character of neural networks means that its information and “knowledge” are not localized in a single component: they are spread out across all the nodes and links, which together contribute to the final solution. This makes the processing much more robust: individual components may be missing, malfunctioning or contain errors; yet, the disturbances this introduces to the process are drowned out by the contributions from the other components when everything is aggregated. In a sequential process, on the other hand,

every step or component through which the process passes constitutes a bottleneck: if that component breaks down, the process may never recover.

The learning happens via a general “reward” or reinforcement mechanism: links that have been successfully used in producing a good solution become stronger; the others become weaker. After many experiences of successful or failed processing, the relative strengths of the different connections will have shifted so that the probability of overall success has become much larger. This intrinsically simple mechanism only works for complex problems because of the distributed character of the processing: if only the process as a whole could be rewarded or punished, this would not produce enough information for it to learn a complex, subtle procedure consisting of many different actions collaborating towards a global solution. Because the process is distributed, its components can learn individually, so that the one can be reinforced at the same time as its neighbor is weakened, thus rebalancing their relative contributions.

Challenges

From problems to opportunities

The view of intelligence as a capability for problem solving or information processing runs into a fundamental issue: what is a meaningful problem, or meaningful information? Why should an intelligent agent address certain problems or process certain information, and disregard others? In other words, how does an agent decide what to do or pay attention to? In the approach of traditional artificial intelligence (AI), this issue is ignored, as AI programs are conceived essentially as question-answering systems: the user or programmer introduces the question (problem, query, input), and the program responds with an answer (solution, output).

On the other hand, the issue becomes inevitable once you start to design autonomous systems, i.e. systems that should be able to act intelligently in the absence of an instructor telling them what to do. Such a system should at least have a *value system*, i.e. a set of explicit or implicit criteria that allow it distinguish “good” outcomes from “bad” ones. Given the ability to evaluate or value phenomena, the agent can then itself decide what aspects of its situation are “problematic” and therefore require some solution.

However, acting autonomously is more than solving problems. A situation does not need to be “bad” in order to make the agent act. When you take a walk, draw something on a piece of paper, or chat with friends, you are not solving the problem of being “walkless”, “drawingless”, or “chatless”. Still, you are following an implicit value system that tells you that it is good to exercise, to play, to be creative, to see things, to build social connections, to hear what others are doing, etc. These kinds of values are positive, in the sense that they make you progress, develop, or “grow” beyond what you have now, albeit without any clear goal or end point. Maslow in his theory of motivation called such values “growth needs” (Maslow, 1955, 1970; Heylighen, 1992). Problems, on the other hand, are defined negatively, as the fact that some aspiration or need is *not* fulfilled. With such “deficiency” needs, once the goal is achieved, the problem is solved, and the motivation to act disappears. This implies a conservative strategy, which is conventionally called “homeostasis”, “regulation”, or “control”: the agent acts merely to compensate perturbations, i.e. phenomena that make it deviate from its ideal or goal state.

The reason that this is not sufficient is evolution: the environment and the agents in it are constantly adapting or evolving. Therefore, no single state can be ideal in all circumstances.

The only way to keep up with these changes (and not lose the competition with other agents) is to constantly adapt, learn, and try to get better. That is why all natural agents have an instinct for learning, development or growth. Therefore, they will act just to exercise, test their skills, or explore new things.

The difference between positive (growth) and negative (deficiency) values corresponds roughly to the difference between positive and negative emotions. Negative emotions (e.g. fear, anger, or sadness) occur when a need is frustrated or threatened, i.e. when the agent encounters a problem that it may not be able to solve. Positive emotions (e.g. joy, love, curiosity) on the other hand, function to broaden your domain of interest and build cognitive, material, or social opportunities or resources (Fredrickson, 2004). In other words, they motivate you to connect, explore, play, seek challenges, learn, experience, etc. Negative emotions tend to narrowly focus your attention to the problem at hand, so that you can invest all your resources in tackling that problem; positive emotions tend to widen your field of attention so that it becomes open to discovering new opportunities for growth.

A general theory of values should encompass both positive or growth values, and negative or deficiency values. From an evolutionary perspective, all values can be derived from the fundamental value of fitness (survival, development, and reproduction), since natural selection has ensured that agents that did not successfully strive for fitness have been eliminated from the scene.

The present paper will assume that intelligent agents have some kind of in-built value system, and assume that those values elicit specific actions in specific situations. For example, in a life-threatening situation, the fundamental value of security or survival will lead the agent to act so as to evade the danger—e.g. by running away from the grizzly bear. On the other hand, in a safe situation with plenty of promise, the value of curiosity will lead the agent to explore a variety of opportunities in order to discover the most interesting ones. The positive or negative intensity of such a situation will be denoted as its *valence*. Valence can be understood as the subjective appreciation by an agent of the global utility, well-being or fitness offered by a particular phenomenon or situation (Colombetti, 2005). It can be represented by a number, which is larger than zero for positive situations, smaller than zero for negative ones, and zero for neutral or indifferent ones.

Definition of challenge

We come to the most important new concept discussed in this paper: a *challenge* is a situation that potentially carries valence for an agent, so that the agent is inclined to act—in the case of negative valence by suppressing the perceived disturbance (s); in the case of positive valence by exploring or exploiting the perceived opportunity (ies). More concisely, we can define a challenge as a *phenomenon that invites action from an agent*.

Negative challenges correspond to what we have called problems; positive challenges represent affordances for growth or progress. But note that these are not opposites but independent dimensions, since a challenge can carry both positive and negative valences. For example, for a hunter, encounter with a wild boar is both an opportunity, since a wild boar has tasty meat, and a problem, since a wild boar is dangerous. For a company, a free trade agreement can be both positive, because it gives access to new clients, and negative, because it opens the door to new competitors. A challenge incites action because it represents a situation in which not acting will lead to an overall lower fitness than acting—because the agent gains fitness by taking action, loses fitness by not taking action, or both. Thus, a challenge can be seen as a *promise of fitness gain for action relative to inaction*.

However, a challenge merely inspires or stimulates action, it does not impose it. The reason is that a complex situation will typically present many challenging phenomena, and the agent will not be able to act on all of them. For example, someone surfing the web typically encounters many pages that seem worth investigating, but obviously cannot read all of them. We may assume that an agent is intrinsically capable of choice, and that this choice will be determined partly by subjective preferences, partly by situational influences, partly by chance, i.e. intrinsically unpredictable, “random” fluctuations. Therefore, it is in general impossible to determine exactly how an agent will react to a situation, although it should be possible to derive statistical regularities about the most common choices.

One of the reasons for this unpredictability is that agents have *bounded rationality* (Gigerenzer & Selten, 2002): they lack the information or cognitive abilities necessary to evaluate all the different challenges. They therefore have to make “informed guesses” about the best course of action to take.

In addition to positivity and negativity, other dimensions worth considering in order to compare challenges are (Heylighen, 2012a):

- *prospect* (in how far can the agent foresee the different aspects or implications of the challenge?),
- *difficulty* (how much effort would be involved in tackling the challenge?), and
- *mystery* (in how far would tackling this challenge increase the agent’s prospect concerning other challenges?).

Prospect distinguishes expected challenges (which direct the agent’s course of action and allow it to work proactively towards (or away from) a remote target) from unexpected ones (which divert the course of action, and force the agent to react). Combining the prospect dimension with different aspects of the valence dimension produces the simple classification of Table 1 (an extension of the one in (Heylighen, 2012a)). The valence dimension has here been subdivided in not only positive, negative and neutral (“indifferent”) values, but also the “unknown” value, which represents the situation where the agent does not (yet) know what valence the challenge may have.

<i>valence</i> \ <i>prospect</i>	Positive	Negative	Unknown	Indifferent
Directions (proactive)	Goals	Anti-goals	Mysteries	Pointers
Diversions (reactive)	Affordances	Disturbances	Surprises	Variations

Table 1: a 2 x 4 classification of challenge types.

Indifferent challenges, while having zero valence, can still function as “challenges” in the sense that they incite actions different from the ones that the agent would take in their absence. For example, a temperature of 15°C, while being neither positive nor negative, requires a different type of clothing than a temperature of 25°C. Indifferent challenges that are foreseen may be called “pointers” or “markers” as they indicate remote phenomena or

circumstances worth taking into account while setting out a course of action. For example, a landmark, such as strangely shaped rock, can help you to orient yourself while walking towards your goal, without being in itself valuable. Indifferent challenges that are not foreseen may be called “variations” or “fluctuations”, as they merely represent the normal type of diversions, such as changes in weather, traffic conditions, people you pass on the street, etc., that are not exactly predictable but not surprising either.

Unknown challenges are potentially much more important than indifferent challenges, as they may turn out to have a high positive or negative valence once more information is gathered. Therefore, they tend to invite action with much more intensity. When their presence is foreseen, they may be called “mysteries” as they represent a focus for curiosity and exploration, inviting the agent to gather additional knowledge. An example would be the entrance to a cave that you can see from afar, however, without knowing what is inside the cave. When they appear unexpectedly, they may be called “surprises” as they function as sudden warnings that the agent’s knowledge has a potentially dangerous gap. An example may be someone shouting at you from across the street, which may be an expression of anger or a greeting.

From activation to relaxation

An advantage of the *challenge* concept is that it is a generalization not only of the *problem* concept, but of the concept of *activation* on which neural network models are built. Indeed, from the definition it follows that a challenge “activates” an agent, by inciting it to act.

In neurophysiology, the more accurate name used to describe neural activation is “action potential”. This denotes a transient rise in the electrical potential of the neuron. This potential is propagated along the neuron’s axon to its outgoing synapses, where it can be transmitted to connected neurons. The underlying mechanism is the following: an increase in potential energy creates a disequilibrium or tension between the parts of neurons that are “activated” and those that are not (that remain at a lower potential).

More generally, in physics a difference in potential energy between two points determines a force that pushes the system from the high potential to the low one. Examples are the voltage that forces electrical current through a wire (or through an axon), or the gravity that pulls a rock down from the hill into the valley. That disequilibrium or force is ultimately what makes the system “active”, what compels it to act. The movement from the higher to the lower potential brings the system back to equilibrium, a process called *relaxation* (Faller, 2012), as it eliminates the tension or potential difference. In the case of a wire or axon, relaxation implies a propagation of the electrical current or activation from the higher to the lower potential.

The same reasoning can be used to understand the resolution of challenges. A challenge can be seen as a *difference* between the present situation (the problem or opportunity) and the ideal situation (the solution to the problem or successful exploitation of the opportunity). Note that the neutral concept of “difference” allows a challenge to be interpreted positively (opportunity) as well as negatively (problem). This difference creates an imbalance or tension that needs to be relaxed, typically by propagating it along some medium until the difference is dissipated. An example is a wave in water or in air: a local disturbance (e.g. a stone thrown into a pond) creates a difference in height or density between the disturbed and non-disturbed parts of the medium; this difference (wave front) then spreads out further until it fades away. In the case of a wave or electrical current, the direction of propagation is obvious: just follow the potential gradient in the direction of steepest descent. In the typical challenges that

confront intelligent agents, the direction is much more complex, as there are many possible routes to increase fitness (i.e. decrease tension). This requires an exploration of different routes, in parallel or in sequence, so as to find the better one. This will bring us to the need to better understand propagation.

An important difference between simple relaxation models and challenge models is that intelligent agents, unlike physical systems, must remain in a far-from-equilibrium state: they are constantly active, consuming energy, and trying to avoid at all costs a complete standstill (i.e. death). Therefore, while they are inclined to relax existing challenges, they will also seek new challenges (affordances, resources, opportunities)—unlike physical systems. In that sense, a “challenge relaxing” dynamics only describes part of their behavior, and must be complemented by a “challenge seeking” dynamics that is better described by some form of active exploration (Heylighen, 2012a). This is the equivalent of what we have called positive or growth values. It is illustrated in the brain by the fact that thinking never stops: activation does not simply diffuse until it fades away; action potentials are continuously generated by the brain itself, even in the absence of outside stimuli that play the role of challenges needing to be relaxed. This may be included in our model by reinterpreting the lack of interesting challenges as a challenge in itself, namely as the problem of boredom. This “metachallenge” can only be relaxed by finding new challenges.

Different agents have in general different value systems, and therefore different “ideal” situations. Therefore, the same situation will produce different challenges for different agents. All agents will try to relax the challenge, i.e. bring it closer to the case where the present situation equals the ideal situation, by acting on it or “processing” it. This allows them to either extract benefit from the opportunity, or avoid the penalty imposed by neglecting the problem. But in general a single agent will not be able to fully exploit an opportunity or fully solve a problem, i.e. completely relax a challenge. This means that the situation after processing by one agent still constitutes a challenge for one or more further agents, who either have a different value system defining the “ideal” situation, or a different set of skills for dealing with the challenge. Thus, some part of the challenge tends to remain, ready to be addressed by other agents.

This produces a complex dynamics of challenge processing and propagation: each agent dealing with a challenge will normally extract some benefit from it, thus relaxing some aspects of the challenge, while leaving some others to be passed on to further agents. If we focus only on the remaining aspects, we see a mechanism of information transmission similar to the spreading of memes (Adar & Adamic, 2005; Heylighen, 1998; Heylighen & Chielens, 2008): messages are communicated from agent to agent, without undergoing much change, until they have reached everyone that may be interested in the message. This could for example describe the diffusion of a particular innovation, fashion, or scientific theory.

If we focus on the challenge aspects that are processed and thus partially relaxed, we see the self-organization of a workflow and division of labor (Busseniers, 2011; Heylighen, 2013): different agents perform different tasks that are part of a common challenge, and then pass on the remaining challenge to others with different skills and/or needs, up to the point where nothing of value is left to extract (i.e. all tasks have been done). To better understand such distributed processing of a challenge we need to investigate the dynamics of propagation.

Propagation of challenges

A generalized concept of propagation

The notion of challenge was introduced as a generalization of the notion of a problem that confronts an individual agent (Heylighen, 2012a). In contrast to the standard paradigm of individual problem solving, the challenge propagation paradigm investigates processes that involve a potentially unlimited number of agents. To deal with this, our initial focus must shift from the agent to the challenge itself: what interests us is how an individual challenge is processed by a collective of agents distributed across some abstract space or network. Instead of an agent traveling (searching) across a space of challenges (problem space), we will consider a challenge traveling (propagating) across a space of agents. This change in perspective is similar to the one that distinguishes memetics from traditional social science models of communication (Heylighen & Chielens, 2008): instead of focusing on the individuals communicating, memetic models focus on the information (“memes”) being communicated.

In general, propagation denotes the spreading or transmission of some recognizable pattern, such as a wave, a species, or an idea. The movement of such a pattern has specific characteristics:

- the interaction is local, as the pattern is initially transmitted only to the immediate neighbors of the point it originated in, who pass it on to their neighbors, and so on...
- the pattern tends to spread outwards so as to cover an ever wider area;
- it tends to change while spreading;
- the pattern requires a physical medium to carry it while propagating;
- this medium has a characteristic topology (such as a 2-dimensional surface for a wave, or a social network for a meme) that affects the shape and extent of the spreading;
- the medium may have additional properties such as time lag, density, or friction that affect the speed of propagation as well as the changes occurring to the pattern.

All these characteristics can be found in messages that are passed along across the Internet, or in activation that spreads across the brain. Since challenges are generalizations of these phenomena, propagation appears like the natural way to describe their dynamics.

Stigmergic and networked propagation

There are two basic cases of challenge propagation: stigmergy and propagation across a network. *Stigmergy* is a mechanism whereby a challenge left by an agent in some medium or workspace that is shared with other agents stimulates those others to further address that challenge (Heylighen, 2015; Parunak, 2006). For example, a paragraph added to a Wikipedia page by one person may incite a second person reading that page to add some extra details, a third one to add a reference for the new material, and a fourth one to correct a grammatical mistake. The reference may then be checked and more accurately formatted by a software agent. Here, challenges are spontaneously addressed by subsequent agents as mediated by the shared space (in this case the Wikipedia page).

In this case of stigmergy, a challenge remains available in a public medium or workspace that all agents can access. If an agent decides to take on the challenge, it will perform some

actions that change the state of the challenge and then leave the modified challenge in the medium. At a later stage, some other agent may pick up the modified challenge, and perform some further work on it, again leaving the “traces” of its work in the medium, where it can function as a challenge for some further agent. The “workflow” from agent to agent self-organizes, as the one leaving the challenge does not know who will pick it up later. Here, the changes in the challenge in a sense propagate in time, but not in space, as they remain in the same place.

In the case of propagation across a social or neural network, the medium has a non-trivial topology that directs the workflow: an agent that has finished working on a challenge passes it on to one or more specific other agents that it is connected to. An example is an email message sent and forwarded with comments from person to person, or a “post” or “tweet” in a social media network that is reposted to other people. Here, a challenge moves from agent to agent by following the available links in the network. In this case, the topology of the network (which node is connected to which other nodes) fundamentally determines the propagation process: a challenge can move directly from agent A to agent B only if there exists a link $A \rightarrow B$ in the network.

In the stigmergic case, the challenge can move from any agent X to any agent Y, without constraints. The only requirement is that Y should “visit” the shared medium some time after X deposited its modified challenge there. In the Wikipedia example, any person can modify any page at any stage independently of which other person has contributed to that page. An example of networked propagation is email, where A can pass a challenge on to B only if A has B’s email address, and B has enough trust in A to take on challenges from A. This typically only happens if A and B have a social or organizational connection.

These are in a sense the extreme cases. What interests us here is the formulation of a more general theory that encompasses both, as well as the ground in between them. An example of such a “middle ground” is an Internet “forum”, i.e. a place where discussions take place between a limited number of people belonging to a specific group or community. All members of the community can post messages (“challenges”) to the forum, read the messages posted by others, and react to those messages (take on the challenge). However, people not belonging to the community can in general not access or create such messages. The forum acts as a private medium for the group. This is similar to stigmergic propagation in that a message is propagated to anyone in the community, but similar to networked propagation in the sense that the message is directed only to members of the community, and to no one else. The Internet as a whole can be conceived as a gigantic collection of such forums, which are partly connected or overlapping, partly disjointed. A forum in the broadest sense can encompass everyone (e.g. anyone can read or write Wikipedia articles), just two people, or anything in between. We will use the term *forum* as the most general form of a “meeting ground” where people can exchange challenges.

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. A complex challenge (say, global warming) has a large number of aspects that require different skills. The problem now is to *distribute* the different challenge aspects across the different agents so as to make sure the challenge as a whole is dealt with efficiently. This is the basic problem of coordination. It includes *division of labor* (who deals with what challenge component?), *workflow* (where does a challenge go after it has been partially dealt with?), and *aggregation* (how are all the finished pieces of work assembled?) (see Fig. 3).

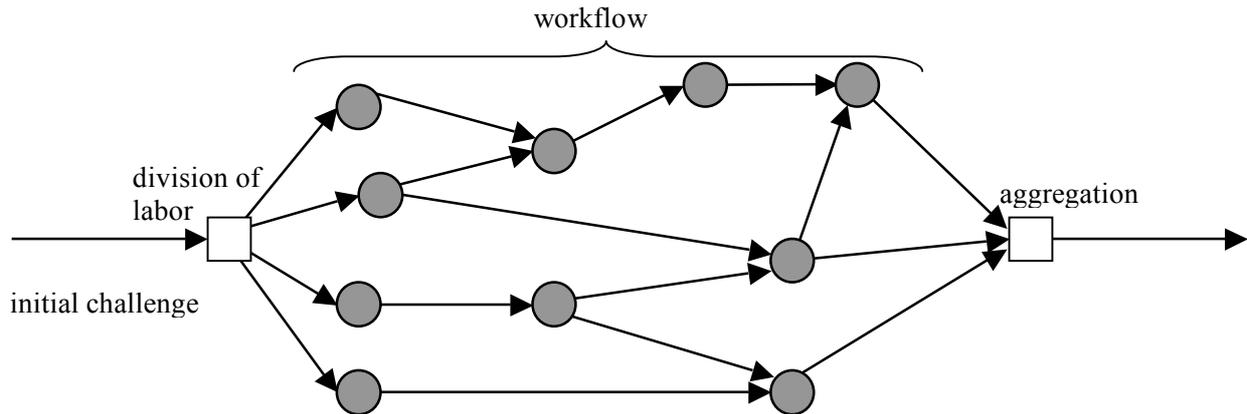


Figure 3: An illustration of coordination, in which an initial challenge is split up in separate activities performed by different agents (division of labor), which are followed by other activities (workflow), and whose results are assembled into a final result (aggregation). Grey circles represent individual agents performing activities. Arrows represent the propagation of challenges from one agent to the next.

Perhaps surprisingly, such distributed coordination can self-organize relatively easily across the Internet, via both stigmergy (Heylighen, 2015) and networked propagation. A good illustration can be found in the different open source communities developing complex software without central supervision (Heylighen, 2007). 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 different components (human and technological) of the network.

Learning in the distributed network

In the case of networked propagation, coordination requires an additional condition, though: the links between the agents that define the network should be appropriate to the task of distributing challenges. Otherwise, challenges are likely to be passed on to agents that do not care about them, or that do not have the appropriate skills to deal with them. Establishing links is achieved via a learning process, which creates and “remembers” adequate links, while “forgetting” inadequate ones.

The similarity between a distributed network of agents and a neural network suggests that the distributed network should be able to learn by differentially strengthening or weakening its links. Delta learning is a form of reinforcement learning (Woergoetter & Porr, 2008) in which a link is rewarded if it brings the challenge closer to relaxation, and penalized if it reduces relaxation. The link strength can then be increased by an amount proportional to the degree of relaxation (which may be negative). The interpretation of this operation is that if an agent

transmits a challenge via a specific link (e.g. sends it to a friend, or posts it to a forum), and it observes that the challenge is adequately dealt with (e.g. the friend provides a good tip on tackling it, or the people on the forum collaboratively develop a solution), then the agent will be more inclined to use the same link in the future to transmit similar challenges. That means that the probability of use of the link, and therefore its weight, increases. Vice-versa, an unsuccessful transmission will decrease the probability of later use.

The network does not need any sophisticated learning mechanisms to adapt in this way to its usage. On the one hand, links strengths will be maintained and updated in people's individual memories as the degree of trust they have in the abilities of others to deal with specific challenges (Van Overwalle & Heylighen, 2006). On the other hand, links will be stored in the external memory that is provided by the worldwide ICT network. For example, links will be created or reinforced by such mundane activities as adding someone's phone or email to your list of contacts, bookmarking a site, linking to someone on a social media network, or registering for some organization (and thus getting easier access to its members and resources). All these activities change the environment of the agent in such a way that this agent becomes more likely to communicate with selected other agents. Moreover, these changes will typically be triggered by successful interactions: you will normally note a person's address if that person was interesting or friendly, join a group if they appear to be doing good work, and bookmark a site if it contains useful information. If later it would turn out that the person, group or site is no longer relevant to your interests, you will similarly weaken your connection with them...

Further Research

The challenge propagation framework as we formulated it here appears like a very promising approach for modeling the complex distributed processes via which problems and opportunities are processed in a self-organizing network. After our conceptual analysis of the main components of the framework, we are ready to define these components and their relationships in a more precise, formal manner. This would not only provide a basis for a mathematical model of challenge propagation, but for a simulation aimed at exploring different variations of the model by investigating how they affect the overall intelligence of the network.

Presently, my research group is developing such a mathematical/simulation model (Heylighen et al., 2012), in order to investigate precisely how the distributed intelligence of the network increases as it selectively strengthens or weakens its links or increases its stigmergic capabilities. Our measure for distributed intelligence is simply the degree to which challenges are resolved by the networked agents as compared to the same group of agents without connecting medium. Our working hypothesis is that distributed intelligence increases as the network learns better connections, and as the number of "forums" for stigmergy increases.

Our preliminary simulation, called *ChallProp* (Veitas, 2012), indeed shows such self-organization of distributed intelligence. However, we will need many more runs with a variety of different parameter settings and variations on the dynamic mechanisms in order to achieve results that are statistically reliable and ready to be applied to more realistic situations. In the meantime, I hope that the present conceptual model will be sufficient to inspire other researchers to apply these ideas in a variety of situations that exhibit distributed intelligence.

References

- Adar, E., & Adamic, L. A. (2005). Tracking information epidemics in blogspace. In *Web Intelligence, 2005. Proceedings. The 2005 IEEE/WIC/ACM International Conference on* (pp. 207–214).
- 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.
- Ben-Jacob, E., Becker, I., Shapira, Y., & Levine, H. (2004). Bacterial linguistic communication and social intelligence. *Trends in Microbiology*, 12 (8), 366–372.
- Busseniers, E. (2011). *Hierarchical organization versus self-organization* (No. 2012-03). Retrieved from <http://pcp.vub.ac.be/ECCO/ECCO-papers/Busseniers-HierOrgvsSelfOrg.pdf>
- Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd, J., Theraula, G., & Bonabeau, E. (2003). *Self-organization in biological systems*. Princeton Univ Pr.
- Clark, A. (1998). Embodied, situated, and distributed cognition. *A Companion to Cognitive Science*, 506–517.
- Colombetti, G. (2005). Appraising valence. *Journal of Consciousness Studies*, 8 (10), 103–126.
- Faller, L. D. (2012). Relaxation phenomenon (physics and chemistry). In *Britannica Online Encyclopedia*. Retrieved from <http://www.britannica.com/EBchecked/topic/496927/relaxation-phenomenon>
- 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
- Fredrickson, B. L. (2004). The broaden-and-build theory of positive emotions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 359 (1449), 1367.
- Gabora, L. M. (1993). Meme and variations: A computational model of cultural evolution. *D. Stein (Ed)*.
- Gigerenzer, G., & Selten, R. (2002). *Bounded rationality: The adaptive toolbox*. the MIT Press.
- Goertzel, B. (2002). *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers.
- Heylighen, F. (1988). Formulating the problem of problem-formulation. *Cybernetics and Systems*, 88, 949–957.
- Heylighen, F. (1992). A cognitive-systemic reconstruction of Maslow's theory of self-actualization. *Behavioral Science*, 37 (1), 39–58. doi:10.1002/bs.3830370105
- Heylighen, F. (1998). What makes a meme successful? Selection criteria for cultural evolution. In *Proc. 16th Int. Congress on Cybernetics* (pp. 423–418). Namur: Association Internationale de Cybernétique.
- 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. doi:10.1023/A:1009690407292

- Heylighen, F. (2007). Why is Open Access Development so Successful? Stigmergic organization and the economics of information. In B. Lutterbeck, M. Baerwolff & R. A. Gehring (Ed.), *Open Source Jahrbuch 2007* (pp. 165–180). Lehmanns Media. Retrieved from <http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf>
- Heylighen, F. (2008). Accelerating socio-technological evolution: from ephemeralization and stigmergy to the global brain. In *Globalization as evolutionary process: modeling global change* (p. 284). Routledge. Retrieved from <http://pcp.vub.ac.be/papers/AcceleratingEvolution.pdf>
- Heylighen, F. (2011a). 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. (2011b). Self-organization of complex, intelligent systems: an action ontology for transdisciplinary integration. *Integral Review*. Retrieved from <http://pespmc1.vub.ac.be/Papers/ECCO-paradigm.pdf>
- Heylighen, F. (2011c). *The GBI Vision: past, present and future context of global brain research* (No. 2011-01). Retrieved from <http://pespmc1.vub.ac.be/papers/GBI-Vision.pdf>
- Heylighen, F. (2012a). *A Tale of Challenge, Adventure and Mystery: towards an agent-based unification of narrative and scientific models of behavior* (Working papers No. 2012-06). Brussels, Belgium. Retrieved from <http://pcp.vub.ac.be/papers/TaleofAdventure.pdf>
- Heylighen, F. (2012b). *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. (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. (2015). Stigmergy as a Universal Coordination Mechanism: components, varieties and applications. In T. Lewis & L. Marsh (Eds.), *Human Stigmergy: Theoretical Developments and New Applications*. Springer. Retrieved from <http://pespmc1.vub.ac.be/papers/stigmergy-varieties.pdf>
- Heylighen, F., & Bollen, J. (2002). Hebbian Algorithms for a Digital Library Recommendation System. In *International Conference on Parallel Processing* (p. 439). Los Alamitos, CA, USA: IEEE Computer Society. doi:10.1109/ICPPW.2002.1039763
- Heylighen, F., Busseniers, E., Veitas, V., Vidal, C., & Weinbaum, D. R. (2012). *Foundations for a Mathematical Model of the Global Brain: architecture, components, and specifications* (GBI Working Papers No. 2012-05). Retrieved from <http://pespmc1.vub.ac.be/papers/TowardsGB-model.pdf>
- Heylighen, F., & Chielens, K. (2008). Evolution of culture, memetics. In R. A. Meyers (Ed.), *Encyclopedia of Complexity and Systems Science* (p. 10370). Springer. Retrieved from <http://pespmc1.vub.ac.be/Papers/Memetics-Springer.pdf>
- Heylighen, F., Heath, M., & Van Overwalle, F. J. (2004). The Emergence of Distributed Cognition: a conceptual framework. In *Proceedings of collective intentionality IV*.

- 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
- Malone, T. W., Laubacher, R., & Dellarocas, C. (2010). The collective intelligence genome. *IEEE Engineering Management Review*, 38 (3), 38.
- Maslow, A. H. (1955). Deficiency motivation and growth motivation.
- Maslow, A. H. (1970). *Motivation and personality* (2nd ed.). New York: Harper & Row.
- McLeod, P., Plunkett, K., & Rolls, E. T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford University Press Oxford.
- Miller, J. H., & Page, S. E. (2007). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, N.J: Princeton University Press.
- Minsky, M. (1988). *The society of mind*. Simon & Schuster.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Prentice-Hall Englewood Cliffs, NJ.
- 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.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. San Diego: University of California Press.
- Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor.
- Trewavas, A. (2003). Aspects of plant intelligence. *Annals of Botany*, 92 (1), 1–20.
- Van Overwalle, F., & Heylighen, F. (2006). Talking nets: A multiagent connectionist approach to communication and trust between individuals. *Psychological Review*, 113 (3), 606.
- Veitas, V. (2012). *Software architecture of the challenge propagation model* (No. 2012-06). Retrieved from <http://pespmc1.vub.ac.be/GBI/vveitas.SoftwareArchitecture.pdf>
- Woergoetter, F., & Porr, B. (2008). Reinforcement learning. *Scholarpedia*, 3 (3), 1448. doi:10.4249/scholarpedia.1448