Challenge Propagation:

towards a mathematical theory of distributed intelligence and the Global Brain

Francis Heylighen
Global Brain Institute
Vrije Universiteit Brussel

1. Introduction

This working paper wishes to introduce a new idea, *challenge propagation*, which synthesizes my older work on spreading activation in collective intelligence [1], and my more recent ontology of action [2]. 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 [3], and complex systems, and which denotes the process by which some phenomenon is iteratively transmitted from a point in a space (or a node in a network) to the neighboring points (or nodes).

The intention of this work is to provide a conceptual and mathematical foundation for a new theory of the Global Brain[4], viewed 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

[1] Francis Heylighen, 'Collective Intelligence and Its Implementation on the Web: Algorithms to Develop a Collective Mental Map', *Computational & Mathematical Organization Theory*, 5 (1999), 253–280 <doi:10.1023/A:1009690407292>; Francis Heylighen and Johan Bollen, 'Hebbian Algorithms for a Digital Library Recommendation System', in *International Conference on Parallel Processing* (Los Alamitos, CA, USA: IEEE Computer Society, 2002), p. 439 <doi:10.1109/ICPPW.2002.1039763>.

[2] Francis Heylighen, 'Self-organization of Complex, Intelligent Systems: An Action Ontology for Transdisciplinary Integration', *Integral Review*, 2011 <http://pespmc1.vub.ac.be/Papers/ECCO-paradigm.pdf>; Francis Heylighen, 'Self-organization in Communicating Groups: The Emergence of Coordination, Shared References and Collective Intelligence', in *Complexity Perspectives on Language, Communication, and Society*, ed. by Angels Massip Bonet and Albert Bastardas (Springer, 2012) <http://pcp.vub.ac.be/Papers/Barcelona-LanguageSO.pdf>.

[3] L. M. Gabora, 'Meme and Variations: A Computational Model of Cultural Evolution', *D. Stein (Ed)*, 1993; Francis Heylighen and K. Chielens, 'Evolution of Culture, Memetics', *Encyclopedia of Complexity and Systems Science* (Springer, 2008), p. 10370 <doi:10.1007/978-0-387-30440-3_189>.

[4] Francis Heylighen, *The GBI Vision: Past, Present and Future Context of Global Brain Research*, GBI Working Papers, 2011; Francis Heylighen, 'Accelerating Socio-technological Evolution: From Ephemeralization and Stigmergy to the Global Brain', in *Globalization as Evolutionary Process: Modeling Global Change*, Rethinking Globalizations, 10 (Routledge, 2008), p. 284; Francis Heylighen, '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.* (Uchitel Publishing, 2011), pp. 274 – 289 <http://pcp.vub.ac.be/papers/GBconceptions.pdf>; B. Goertzel, *Creating Internet Intelligence: Wild Computing, Distributed Digital Consciousness, and the Emerging Global Brain* (Kluwer Academic/Plenum Publishers, 2002).

network researchers have shown is distributed over the millions of neurons in the brain [5]—, the collective intelligence of insects, but also various as yet poorly understood forms of intelligence in e.g. bacteria [6] or plants [7].

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" [8]: a mind or intelligence can be seen a collaboration between relatively independent modules or 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 [9], 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. Yet, there exist already a number of successful paradigms for doing this, including multi-agent systems, complex dynamic systems, neural networks, and stigmergy[10]. The challenge propagation paradigm is intended to synthesize the best features of these different models. The present paper will outline the basic components that are necessary to build such an integrated mathematical and computational model, including the questions that still need to be resolved before such a model can be implemented.

2. A brief review of intelligence models

The most simple and common definition of intelligence is *the ability to solve problems* [11]. 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)

[5] W. Bechtel and A. Abrahamsen, *Connectionism and the Mind: An Introduction to Parallel Processing in Networks*. (Basil Blackwell, 1991); P. McLeod, K. Plunkett and E. T Rolls, *Introduction to Connectionist Modelling of Cognitive Processes* (Oxford University Press Oxford, 1998).

[6] E. Ben-Jacob and others, 'Bacterial Linguistic Communication and Social Intelligence', *Trends in Microbiology*, 12 (2004), 366–372.

[7] A. Trewavas, 'Aspects of Plant Intelligence', *Annals of Botany*, 92 (2003), 1–20.

[8] Marvin Minsky, *The Society of Mind* (Simon & Schuster, 1988).

[9] Heylighen, 'Self-organization in Communicating Groups'.

[10] Francis Heylighen, *Stigmergy as a Generic Mechanism for Coordination: Definition, Varieties and Aspects*, Ecco, 2011 <http://pespmc1.vub.ac.be/Papers/Stigmergy-WorkingPaper.pdf>; H. Van Dyke Parunak, 'A Survey of Environments and Mechanisms for Human-human Stigmergy', in *Environments for Multi-Agent Systems II* (Springer, 2006), pp. 163–186.

[11] Heylighen, 253–280.

[12]. This requires determining the right sequence of steps that leads from *x* to *y*. 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. This is one trillion for a path of merely 12 steps! 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”.

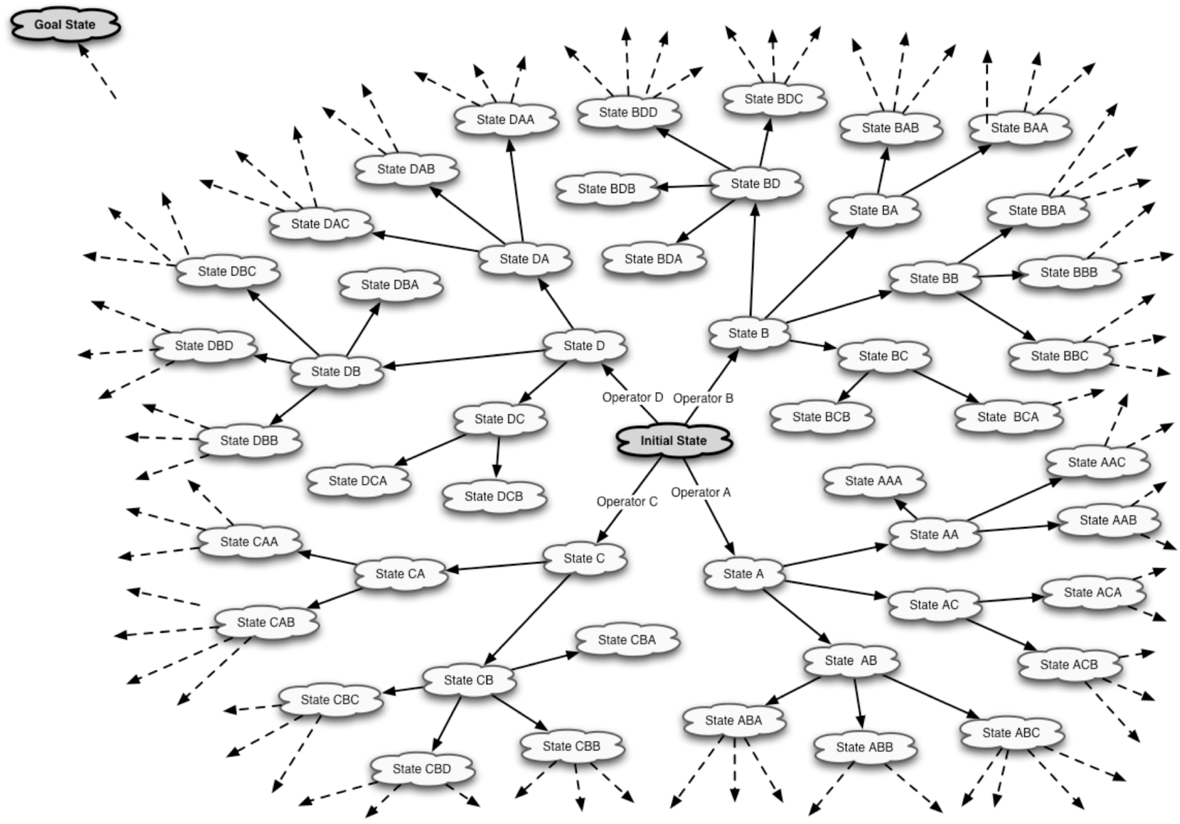

The more problems an agent can solve, the more intelligent it is. Note that this definition does not provide an absolute measure of intelligence, as the number of problems that a non-trivial agent can solve is typically infinite. Therefore, counting the number of solvable problems does not produce the 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.

The partial order provides us with an unambiguous criterion for 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

[12] Francis Heylighen, ‘Formulating the Problem of Problem-formulation’, *Cybernetics and Systems*, 88 (1988), 949–957; A. Newell and H. A Simon, *Human Problem Solving* (Prentice-Hall Englewood Cliffs, NJ, 1972).

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 most significant problems in your present environment. Adding the capability to solve some purely theoretical problems that are of no use 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 x 7, then the solution is 49. The mathematical structure of this model is simply a function that maps every problem state x onto the corresponding solution state y:

$$f: x \rightarrow f(x) = y.$$

The next, more complex model of intelligence is a deterministic *algorithm*. This is a fixed-length or iterated sequence of actions that need to be performed on the initial state, until the state they produce satisfies the condition for being a solution. An example is a procedure to calculate 734 x 2843 or a program that determines the first 100 prime numbers. Such deterministic procedures to manipulate numbers, or more generally, strings of symbols, have given rise to the notion of intelligence as *computation*.

A basic algorithm is guaranteed to produce a 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 generates 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 steps are performed one by one. A first step in our intended generalization towards distributed, parallel 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 very common, its meaning remains surprisingly vague: how exactly is a given piece of information transformed into a new—presumably more useful or meaningful—piece of information? Apart from algorithmic 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 [13]. 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

[13] McLeod, Plunkett and Rolls.

or more "hidden" layers, to the final "output layer". 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 properties 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.

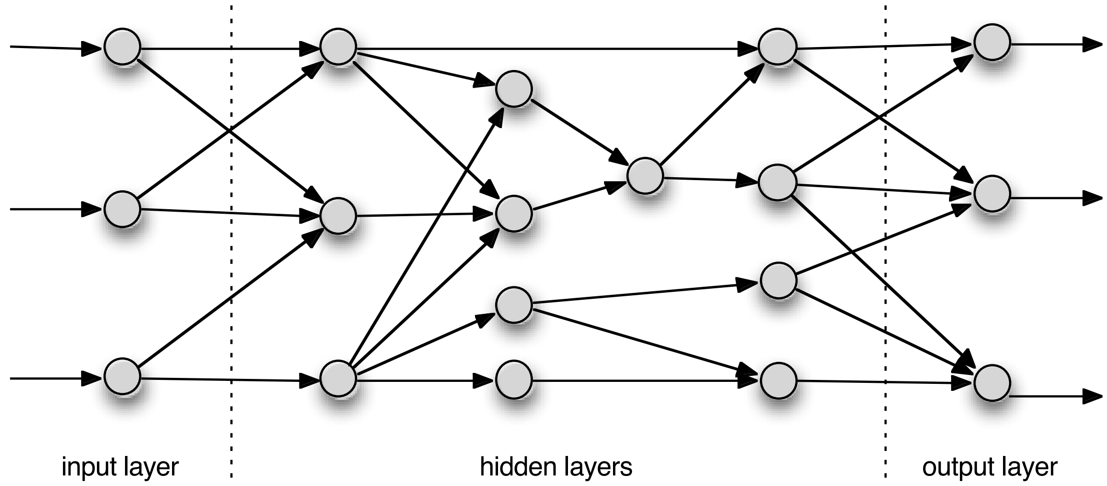

The more general version of such a "feedforward" network is called a "recurrent" network. The difference is that a recurrent network allows activation to cycle back to nodes activated before. Thus, there is no imposed direction "forward", from input layer to output layer. The input in this case is simply 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:

1) processing happens in a parallel, distributed manner, making it more robust and flexible;
2) the network does not need an explicit program to know how it should function: 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.

3. Challenges

3.1 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" [14]. 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 roughly corresponds 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 perturbation that it may not be able to compensate. Positive emotions (e.g. joy,

---

[14] Abraham H. Maslow, 'Deficiency Motivation and Growth Motivation.', 1955; Abraham H. Maslow, *Motivation and Personality*, 2nd edn (New York: Harper & Row, 1970); Francis Heylighen, 'A Cognitive-systemic Reconstruction of Maslow's Theory of Self-actualization', *Behavioral Science*, 37 (1992), 39–58 <doi:10.1002/bs.3830370105>.

love, curiosity) on the other hand, function to broaden your domain of interest and build up cognitive, material, or social resources [15]. 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. The present paper will not further develop such a theory. Yet, it is worth pointing out that such a theory would be an important contribution to a general model of intelligence and of the global brain, and therefore definitely worth investigating. Some inspiration can be found in the various psychological theories of motivation or needs, which include fundamental needs/values such as security, social affiliation, achievement and knowledge (see e.g. [16] [17]). More generally, 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[18]. It can be formalized as a scalar variable, which is larger than zero for positive situations, smaller than zero for negative ones, and zero for neutral or indifferent ones.

3.2 Definition of challenge

We come to the most important new concept discussed in this paper: a *challenge* is a situation that 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, since it gives access to new clients, and negative, since 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

[15] B. L Fredrickson, 'The Broaden-and-build Theory of Positive Emotions.', *Philosophical Transactions of the Royal Society B: Biological Sciences*, 359 (2004), 1367.
[16] Maslow, *Motivation and Personality*.
[17] Heylighen, 39–58.
[18] G. Colombetti, 'Appraising Valence', *Journal of Consciousness Studies*, 8 (2005), 103–126 <http://polorovereto.unitn.it/%7Ecolombetti/docs/GC_AppraisingValence05.pdf>.

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. The implication for modeling is that an agent should not be represented as a deterministic automaton, but as a stochastic system, which may make different decisions in apparently identical cases, but where it is meaningful to specify conditional probabilities for the different choices, so that in a given condition e.g. 50% of agents can be expected to make choice A, 30% choice B, and 20% choice C.

One of the reasons for this unpredictability is that agents have *bounded rationality*[19]: they normally never have all the information or cognitive abilities necessary to evaluate the different challenges. They, therefore have to make “informed guesses” about the best course of action to take. Moreover, assuming that similar agents tend to look for similar resources, it is worth making a choice different from the choice of the others, so as to avoid competition for scarce resources.

In addition to positivity and negativity, other dimensions worth considering in order to compare challenges are:

- *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[20]). 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.

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

[19] G. Gigerenzer and R. Selten, *Bounded Rationality: The Adaptive Toolbox* (the MIT Press, 2002).
[20] Francis Heylighen, *A Tale of Challenge, Adventure and Mystery: Towards an Agent-based Unification of Narrative and Scientific Models of Behavior*, ECCO Working Papers, 2012.

Table1: 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 different actions than 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 that may be taken into account while setting out a course of actions. 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, chance encounters, 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 functions as sudden warnings that the agent's knowledge has a potentially dangerous gap. An example would be a hole that suddenly opens up in the ground before your feet.

3.3 Vector representation of challenges

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 stimulating it to act. The generalization is that challenges are typically complex and multidimensional, just like problems, while activation is normally a one-dimensional quantity (typically varying between 0 and 1). On the other hand, challenges can be fuzzy and vary continuously—like activation, but unlike traditional problems.

The simplest way to formalize this complex nature is to represent challenges as vectors, i.e. points in a multidimensional state space, characterized by values for each of the relevant variables or dimensions. Such a vector is just a list of numbers, e.g. (0.37, 2.4, 7.23, …). A possible simplified representation is a string of binary digits, e.g. (0, 1, 1, 1, 0, 0, …). For flexibility of modeling, it may be worth adopting the trick used in classifier systems [21], and allow such lists to contain *unspecified* numbers, i.e. variables without a value attached to them. Examples are (0.37, #, 7.23, …) where # can be any number, or (0, 1, 1, #, 0, #, …) where # can be either 1 or 0. This means that the agent does not know the value of the corresponding variable, or that the value is not determined (yet). This allows us to represent indifferent or unknown challenges.

A vector representation is used implicitly in neural network models since the activation of the input layer of a network defines an activation vector, with each node representing an independent activation variable. Therefore, a possible implementation of challenge propagation could start from autonomous agents that each have a neural network for individual processing, and that communicate by sending activation vectors (rather than

[21] John H. Holland and others, *Induction: Processes of Inference, Learning, and Discovery* (MIT Press, 1989).

one-dimensional activation values) to the agents they are connected with, as used in our *Talking Nets* model[22]. Another implementation may build further on the classifier system formalism [23], where binary strings are "posted" on a "message board", where they can be picked up by other agents that somehow recognize (part of) the string as relevant to their own interests.

3.4 From activation to relaxation

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, which is propagated along its 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* [24], 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 initial state) and the ideally reachable situation (the goal or the opportunity). Note that the neutral concept of "difference" allows 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 ever further until it completely fades away. In the case of waves or electricity, 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 (=decrease tension), and most routes end in local optima that are less good than the global optimum. This requires a process of exploration of different routes, in parallel or in sequence, so as to find the better one. This brings us to the need to model propagation.

A difference between simple relaxation models and challenge models is that intelligent agents (as living or artificial systems), 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. This is the equivalent of

---

[22] Frank Van Overwalle and Francis Heylighen, 'Talking Nets: A Multiagent Connectionist Approach to Communication and Trust Between Individuals.', *Psychological Review*, 113 (2006), 606.
[23] Holland and others.
[24] Larry D. Faller, 'Relaxation Phenomenon (physics and Chemistry)', *Britannica Online Encyclopedia*, 2012 <http://www.britannica.com/EBchecked/topic/496927/relaxation-phenomenon> [accessed 16 January 2012].

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.

The "difference" interpretation of challenges can be easily formalized using the vector representation. The vector **c** representing a challenge can be decomposed as the difference between the actual situation vector **s** and a vector **g** representing the ideal situation or goal for a specific agent a_i:

$$\mathbf{c}_i = \mathbf{s} - \mathbf{g}_i$$

We may assume that different agents are characterized by different ideals that depend on their value system. Therefore, the same situation **s** will produce different challenges for different agents. All agents will try to relax the challenge, i.e. reduce it to the nil vector: $\mathbf{0} = (0, 0, 0, \ldots)$, which represents the case where the present situation equals the ideal situation.

The challenge seeking dynamics can be represented by the fact that the agents are constantly confronted by a variety of new challenges, and that they are motivated to pick out the most challenging one of those as soon as the previous challenge has been relaxed. These challenges have an external origin: they are produced by the agent's environment (which includes other agents as well as the natural and technological environment). We may assume that this environment is in a constant flux, so that agents are "showered" with challenging phenomena. Each of those is an opportunity to extract valence by acting.

The flux of challenges is similar to the flow of energy that passes through a far-from-equilibrium system (such as a living organism): the system can maintain itself only by extracting (and eventually dissipating) energy from that flow. Typically, non-equilibrium systems tend to self-organize into what Prigogine called "dissipative structures", so as to maximally dissipate this incoming energy, thus generating a maximum of thermodynamic entropy[25]. Similarly, we may assume that agents and systems of agents will tend to self-organize so as to maximally extract benefit from the incoming stream of challenges, because that is what maximizes their fitness in the evolutionary sense. But to do that, they need an efficient mechanism to select the challenges that promise the largest benefit.

3.5 The need for an attention mechanism

Challenges have been defined as "invitations" to act, which means that the agent is not forced to take up the invitation. This is necessary because in general many challenges vie for an agent's attention, and not all can be taken up. This becomes particularly clear if we imagine a hunter-gatherer exploring a stretch of rain forest with its thousands of varied stimuli, or a person surfing the web with it millions of potentially interesting pages. The agent's capabilities for cognitive processing and action are limited, and therefore a selection must be made between these myriads of potentially important challenges. A model of challenge propagation ideally should include a system of selection criteria and mechanisms that help the agent single out the most significant challenges—however, without assuming that there is a single most important one.

---

[25] G. Nicolis and I. Prigogine, *Self-organization in Nonequilibrium Systems: From Dissipative Structures to Order Through Fluctuations* (Wiley, New York, 1977); I. Prigogine and I. Stengers, *Order Out of Chaos: Man's New Dialogue with Nature* (Bantam books, 1984).

Some work that may be of help here is Baars's model of consciousness [26], which assumes that within the brain many stimuli and thoughts compete for attention. However, only a single one can be amplified to the degree that it comes to dominate the "global workspace" that interconnects all the more specialized brain regions. From this workspace, the thoughts selected to be most important for conscious attention are "broadcasted" to the whole brain. This model presupposes a positive feedback mechanism where strongly activated thoughts become even stronger, while suppressing weaker ones, until a single "winner" remains. This is an example of the "winner-takes-it-all" dynamic common in non-linear, self-organizing systems. The suppression requires a mechanism of neural inhibition that may use some form of negative activation [27].

Research on attention and consciousness (see my lecture notes on "Cognitive Systems"[28] for a review) points to unexpectedness and valence as major criteria for winning the competition: stimuli that are surprising [29], and/or that are highly relevant for the present goal tend to attract most attention. The valence criterion is most obvious, as challenges that promise a higher reward than others deserve more attention.

The surprise criterion functions to redirect our attention whenever something unexpected happens, so that we can quickly ascertain whether this novel phenomenon is a danger or an opportunity. In Table 1, we defined a *surprise* as a *diversion* (a challenge that was not predicted, i.e. for which there was no prospect) with unknown ("mysterious") valence. Thus, a surprise is to be distinguished from other diversions (e.g. a gust of wind), whose valence (positive, negative or zero) is known.

Challenge selection is similar to memetic selection, i.e. the selection of ideas ("memes") that are interesting enough to be communicated to others. In earlier work, I have developed a list of memetic selection criteria[30], which overlaps with the simpler SUCCES model [31]. These criteria can be grouped into the ones that determine how easily an idea will be assimilated by an individual (objective and subjective criteria) and those that determine how easily it will be propagated to others (intersubjective criteria). The full list of criteria is too long to review here, but a general idea of the individual criteria can be gotten from the SUCCES acronym, which stands for: *Simple, Unexpected, Concrete, Credible, Emotional Stories*. We already mentioned two of these: *surprise* (unexpectedness and mystery), and *valence* (which is the basic dimension distinguishing positive from negative emotions).

Concreteness is a criterion that was formalized in the classifier system model of cognition [32], where "messages" (challenges, ideas) posted on the cognitive system's "message board" (workspace, propagation medium) are selected in part on the basis of their degree of specificity: more general messages are processed with a lower priority than more specific ones. "Generality" in this formalism is measured simply as the number of # symbols ("unspecified") in the list of values: the smaller this number, the higher the priority of the challenge.

---

[26] B.J. Baars, 'The Global Workspace Theory of Consciousness', *The Blackwell Companion to Consciousness*, 2007, 236–246.
[27] S. Dehaene, M. Kerszberg and J.P. Changeux, 'A Neuronal Model of a Global Workspace in Effortful Cognitive Tasks', *Proceedings of the National Academy of Sciences*, 95 (1998), 14529.
[28] Francis Heylighen, *Cognitive Systems*.
[29] Jeff Hawkins and Sandra Blakeslee, *On Intelligence* (Henry Holt and Company, 2005).
[30] Heylighen and Chielens, p. 10370; Francis Heylighen, 'What Makes a Meme Successful? Selection Criteria for Cultural Evolution', in *Proc. 16th Int. Congress on Cybernetics* (Namur: Association Internat. de Cybernétique, 1998), p. 423–418.
[31] C. Heath and D. Heath, *Made to Stick: Why Some Ideas Survive and Others Die* (Random House, 2007).
[32] Holland and others.

A practical application of this criterion can be found in the GTD (*Getting Things Done*) methodology [33]: "next actions" (tasks, challenges) that are formulated more concretely are more likely to be executed quickly. For example, a task formulated as "arrange meeting" is much less specific than one formulated as "call John to ask whether Thursday is OK for meeting". Someone who is hesitating about which of several tasks to do now will be much more inclined to pick up the second task, since this task does not require further reflection or information gathering about what exactly needs to be done.

The "story" criterion notes that ideas are more likely to be assimilated if they are presented in the form of a *narrative*, i.e. an account of a course of action performed by some agent. The reason is that the person listening to the story tends to empathize with the "hero" of the story, that is, imagine performing the actions herself or himself [34]. This internal simulation of the course of action mentally prepares the person to perform similar actions in order to tackle similar challenges. This may be seen as another aspect of the concreteness criterion: the more salient, intuitive and specific the suggested action, the easier it is for the agent to execute it, and therefore the more the agent will be inclined to effectively start acting on the challenge.

Credibility is simple a measure of how trustworthy or reliable the information is. When you are uncertain whether a challenge really exists, you will be less inclined to tackle it. For example, a warning coming from a weird religious cult that "a comet will strike the Earth" is unlikely to incite you to act, while you would take one coming from NASA much more seriously. Credibility too can be seen as an addition to the concreteness criterion, since it reduces the agent's uncertainty about what to do next. The same can be said about the *simplicity* criterion, since simple information is easier to process and interpret, and therefore leaves less doubt about how to deal with it.

In sum, these criteria can all be seen as aspects of what we might call the *clarity* dimension of challenges: the clearer the view the agent has of the challenge, the lower its uncertainty about how to act, the more the agent will be inclined to effectively act. A possible formalization of clarity may be to attach a probability P to each of the possibilities i for action suggested by a challenge, and then calculate the uncertainty (statistical entropy H) of that probability distribution according to the classic Shannon formula:

$$H = -\sum_i P_i . \log P_i$$

Higher uncertainty then corresponds to lower clarity, with maximal clarity corresponding to zero uncertainty. This maximum clarity is what we find in a look-up table model of intelligence: when the challenge is "what is 7 x 7?", then the right action, with absolute certainty, is "write down 49". But this kind of intelligence is trivial, of course, and true intelligence only starts to shine its light when the situation is intrinsically less clear…

We are left with three fundamental criteria for deciding which challenges are most worth attending to: *valence*, *surprise*, and *clarity*. An agent is most likely to act on a challenge that entails high value (large potential gains or losses), that is mysterious, unusual or surprising, and where there is minimal uncertainty or ambiguity about how to act. Note that the surprise/mystery and clarity criteria at first sight may appear to be opposite. However, the uncertainty implied by surprise is one about the nature of the challenge (what is its true valence?), while the certainty implied by clarity is one about the choice of action (how should I deal with this challenge?). These are not contradictory, as an uncertain challenge (e.g. a

---

[33] D. Allen, *Getting Things Done: The Art of Stress-free Productivity* (Penguin Group USA, 2001); Francis Heylighen and Clément Vidal, 'Getting Things Done: The Science Behind Stress-Free Productivity', *Long Range Planning*, 41 (2008), 585–605 <doi:10.1016/j.lrp.2008.09.004>.

[34] C. Heath and D. Heath; Heylighen, *A Tale of Challenge, Adventure and Mystery: Towards an Agent-based Unification of Narrative and Scientific Models of Behavior*.

mysterious envelope left on your doorstep) may well incite a clear action (e.g. open the envelope to see what it contains). It is clear that in such a case the challenge will have a high priority even though its valence is as yet unknown: you are likely to immediately open the envelope, even though you may be disappointed to see that it contains merely some irrelevant publicity.

3.6 Extracting benefit from challenges

There are in essence two ways in which agents can get benefit from a challenge: either they "consume" the benefit, so that nothing is left, or they merely "make use" of it, so that it remains available for other agents. The first case characterizes material or energetic resources, which the total amount of resources is conserved: whatever one agent takes is no longer available for others. This is the essence of a "zero-sum game": whatever you gain, I lose (and vice versa). For example, the food you eat can no longer be eaten by me. Such resources are called "rival" in economics. "Non-rival" resources do not obey such conservation laws: your gain does not prevent me from gaining an equivalent amount as well. This is a positive-sum game. This is typical for informational resources: when I give you a piece of information, I haven't lost anything, since I still have the same information in my memory[35].

We may assume that challenge vectors contain two types of components: those representing rival phenomena, and those representing non-rival ones. When an agent tackles a challenge it will typical subtract from the rival components, but leave the non-rival ones in place (while still extracting benefit from them). That means that after an agent has dealt with a challenging situation, a new situation will remain that in general constitutes a challenge for one or more agents "further down the line". If agents were to maximally reduce all components of a challenge, only the $\mathbf{0}$ vector would be left, implying that no one else could get any benefit from this challenge anymore. The reason this rarely happens is threefold:

1) as noted, non-rival components will not be reduced. This means that purely informational challenges (e.g. announcements, songs, memes…) can be propagated indefinitely without losing their value;
2) rival components will only be reduced insofar that the agent has the specific skill or ability to tackle these components. If it does not, this leaves potential benefit to be extracted by other agents with different skills. This is the basis for the division of labor: let others do what you cannot do;
3) different agents have in general different need vectors $\mathbf{g}_i$; therefore, the same situation $\mathbf{s}$ will be interpreted by them as different challenges $\mathbf{c}_i = \mathbf{s} - \mathbf{g}_i$, so that e.g. $\mathbf{c}_1 = \mathbf{0}$, but $\mathbf{c}_2 \neq \mathbf{0}$. This means that the "waste product" of one agent (whatever is left after it has extracted everything that it considers beneficial) may still provide a resource for another agent with different needs. This is the basis for complementarity: associate with others that have complementary needs.

These different mechanisms produce a complex dynamics of challenge processing and propagation: each agent dealing with a challenge will normally extract some benefit from it, while reducing some components, and leaving some others invariant. If we focus only on the invariant components, we get a model of information transmission similar to the spreading of

[35] Francis Heylighen, 'Why Is Open Access Development so Successful? Stigmergic Organization and the Economics of Information', in *Open Source Jahrbuch 2007*, ed. by B. Lutterbeck, M. Baerwolff & R. A. Gehring (Lehmanns Media, 2007), pp. 165–180 <http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf>.

memes[36]: messages are passed on from agent to agent, without undergoing much change, until they have reached everyone that may be interested in the message. This could for example be a model of the diffusion of a particular innovation, fashion, or scientific theory.

If we focus on the variable components, we get a model of the self-organization of workflow and division of labor[37]: 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 will need to investigate the dynamics of propagation from agent to agent.

4. Structure and dynamics of the propagation medium

4.1 A generalized concept of propagation

In general, propagation denotes the spreading or transmission of some recognizable pattern, such as a wave, a species of plant, or an idea. What interests us here is that the movement of such a pattern has specific characteristics:

1) 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…
2) the pattern tends to spread outwards so as to cover an ever larger area;
3) it tends to change while spreading, like a wave diminishing in amplitude;
4) some part or trace of the pattern may remain in the places through which it passed;
5) the pattern needs a physical medium to carry it while propagating;
6) 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;
7) 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 information that gets 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.

The notion of challenge was initially introduced as a generalization of the notion of a problem or opportunity that confronts an individual agent[38]. That agent typically deals with the challenge by following a particular course of action, e.g. it searches for a solution to the problem, until the "tension" (difference between situation and need) is "relaxed" (minimized). The standard way to formalize such a process consist in defining a state space, where the states represent possible variations of the challenge or problem situation, and then determining a trajectory that the agent will follow through that space, i.e. a sequence of neighboring states starting at the initial challenge and ending in the "solution" of the problem. The dynamics of the system is typically represented by one of the following formalisms:

- a fitness function on the state space, transforming it into a fitness landscape. Here the agent typically chooses the trajectory characterized by the steepest ascent, i.e. that makes it climb as quickly as possible to a fitness peak
- a vector field or gradient on the state space, which points from each state to the subsequent state. This can be represented graphically as a "phase portrait", i.e. a space

[36] Heylighen, p. 423–418; Heylighen and Chielens, p. 10370; E. Adar and L. A Adamic, 'Tracking Information Epidemics in Blogspace', in *Web Intelligence, 2005. Proceedings. The 2005 IEEE/WIC/ACM International Conference On*, 2005, pp. 207–214.
[37] Heylighen, 'Self-organization in Communicating Groups'; Eva Busseniers, *Hierarchical Organization Versus Self-organization*, GBI Working Paper, 2012.
[38] Heylighen, *A Tale of Challenge, Adventure and Mystery: Towards an Agent-based Unification of Narrative and Scientific Models of Behavior*.

filled with arrows, where the arrows indicate the direction of movement to the next state. Mathematically, such dynamics can be represented by a difference equation, which expresses the next state as a function of the present state: $s(t+1) = f(s(t))$

- a Markov process, where the transition from any state s_i to any other state s_j has a fixed probability $0 \leq P(s_i \rightarrow s_j) \leq 1$ thus defining a transition matrix

In contrast to this standard paradigm for modeling individual problem solving, the challenge propagation paradigm models processes involving a potentially unlimited number of agents. To deal with this, our initial focus shifts from agent to challenge: what interests us is how an individual challenge is processed by a collective of agents distributed across some abstract space or topology. Instead of an agent traveling (searching) across a space of challenges, 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[39]: instead of focusing on the individuals communicating, memetic models focus on the information ("memes") being communicated.

Mathematically, the agent-centered and challenge-centered representations are to some extent a dual or complement of each other: while the one turns the other one upside-down, the two representations are largely isomorphic as they both consider entities of one type (agents or challenges) encountering a sequence of entities of the other type (challenges or agents), while traveling across a space in which the second type of entities are distributed according to a particular topology. This isomorphism or duality is imperfect, though: agents typically deal with only a single challenge at a time, while challenges can reach several agents simultaneously. The latter "spreading" property is what enables parallel or distributed processing, with its many advantages mentioned earlier. The price to pay for this increased processing power is a more complex mathematical and computational model. An analogy from physics may be the increased complexity of field theories or wave models compared to particle models.

However, a challenge propagation model still has the advantage over a wave/field model that agents are finite in number and distributed across a discrete network rather than smeared out across an infinite, continuous space. This discrete character makes computer simulation much more tractable (while making analytic solutions less realistic). Another difference between agent and challenge models is that agents normally maintain a stable identity[40] while traveling, while challenges change up to the point where they may vanish completely ("relaxation" of the challenge).

4.2 Topology of the propagation medium

There are two paradigmatic cases of challenge propagation: stigmergy and propagation across a network. In the case of stigmergy [41], 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 leaves 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. Here, the changes in the challenge in a sense propagate in time, but not in space, as they always remain available in

[39] Heylighen and Chielens, p. 10370.
[40] Heylighen, 'Self-organization of Complex, Intelligent Systems'.
[41] Heylighen, *Stigmergy as a Generic Mechanism for Coordination*; Parunak, pp. 163–186.

the same place. The medium either has a trivial topology (there is only one "place"), or its topology does not directly affect the propagation from agent to agent[42].

The other paradigmatic case is a social or neural network, in which a challenge moves from agent to agent by following the available links in the network. Here, 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 → 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. An Internet example of stigmergy is Wikipedia, where any person can modify any page at any stage independently of which other person has contributed to that page. An Internet 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—which typically only happens if A and B have a social or organizational connection.

These two cases are extremes, and what interests us here is the formulation of a more general formalism that encompass both, as well as the ground in between them. This "middle ground" can be illustrated by 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 generally cannot 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 communication in the sense that the message is directed only to members of the community, to no one else. The Internet as a whole can be conceived as a gigantic collection of such forums, which are partly overlapping, partly disjointed. A forum in this broad 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.

One way to formalize this view of the Internet is by using the mathematical concept of a *hypergraph*. A traditional graph or network is defined as a set N of *nodes* (also called "vertices"), together with a set L of *links* (also called "edges" or "arcs"). The links are defined as pairs of nodes:

$$l \in L, l = (n_1, n_2) \text{ with } n_1, n_2 \in N.$$

In a hypergraph the corresponding "hyperedges" h (which we will call *multilinks* to avoid confusion with the notion of "hyperlink" used in hypertext) are defined as sets of nodes:

$$H = \{h_1, h_2, h_3, \ldots\}, \text{ with } h_i = \{n_1, n_2, n_3, \ldots\}, n_i \in N.$$

This is the definition for an undirected hypergraph, i.e. one in which the links work in both directions: if n_i is linked to n_j, then n_j is necessarily linked to n_i. The more general hypergraph is *directed*[43]: a link in one direction does not necessarily imply the existence of a link in the opposite direction. Such a directed multilink can be expressed as a pair of two sets of nodes, the set of input nodes H_{in} and the set of output nodes H_{out}. The interpretation is that a link

[42] Parunak, pp. 163–186.

[43] G. Gallo and M. G Scutella, 'Directed Hypergraphs as a Modelling Paradigm', *Decisions in Economics and Finance*, 21 (1998), 97–123; G. Gallo and others, 'Directed Hypergraphs and Applications', *Discrete Applied Mathematics*, 42 (1993), 177–201.

exists from any node in the input set to any node in the output set. If the two sets are identical, the corresponding multilink is symmetric or undirected, as it does not distinguish between input and output.

In our model, a multilink can be seen as an intermediary or conduit for the propagation of challenges between linked agents. The input set consists of agents that can deposit challenges into the multilink. These can be then be picked up by the agents that form the output set. On the Internet, a typical multilink is a mailing list or forum to which only a selected group of people may post messages, while a selected group of people can read those messages. In the limit case of full stigmergy, that group is everybody on the planet. In the limit of personal email, the input group contains just the single sender, and the output group the single receiver of the message.

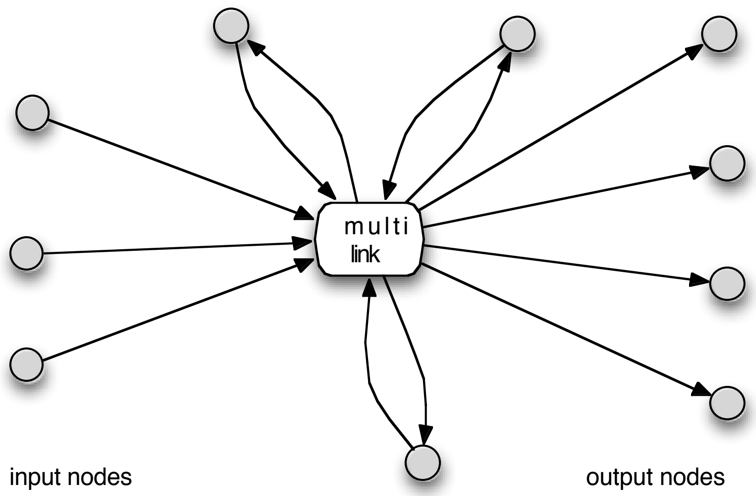

Figure: a multilink in a hypergraph can be seen as an intermediary connecting a set of input nodes with a set of output nodes. Note that some nodes can belong both to the input and to the output of the multilink.

Note that a hypergraph can always be transformed into a traditional, "binary" graph by representing the multilinks as a new type of node (see Fig.). The original nodes then have one or more incoming, binary links to one or more "multilink" nodes, which themselves have outgoing, binary links to ordinary nodes. The transformation just creates two types of nodes with similar properties, but with the constraint that two nodes of the same type cannot be directly linked to each other: there must always be a node of the other type in between.

In the challenge propagation paradigm, the multilink nodes correspond to a special type of "conduit" agents that do not process challenges, but merely store them, perhaps aggregate several of them, and pass them on to ordinary, "active" agents. A conduit agent could be a website, a forum, a book, a Facebook page, a mailing list, or an email connection between two individuals.

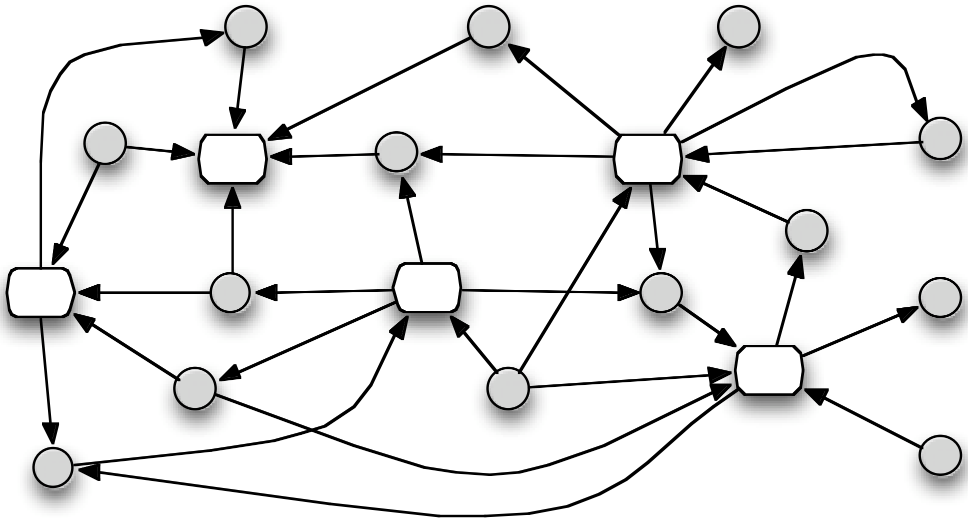

Figure: A directed hypergraph represented as an ordinary directed graph with two types of nodes: ordinary nodes (grey circles) and "multilink nodes" (white ovals). The first can be interpreted as agents, the second as intermediaries along which challenges are propagated from agent to agent. Note that only nodes of different types can be directly linked.

Given this general sketch of the "hyper" network connecting agents, we can now investigate different types of networks and see how they affect the propagation of challenges. Some classic varieties of networks[44] are directed vs. undirected networks, recurrent vs. "feedforward" networks, random networks, "small world" networks, clustering networks, and networks with a power-law distribution for the number of links per node. We will also need to explore variants that are specific to hypergraphs, e.g. by varying the number of nodes that connect to a multilink. As a general rule of thumb, the most interesting networks to explore will be those that are most similar in structure to the "natural" networks that we find in social networks, the Internet, or the brain. These all tend to have the following properties:

- recurrent (links are directed but can form loops or cycles)
- small-world (two arbitrary nodes tend to be separated by not more than a small number of links)
- clustering (two neighbors of a node tend to be connected to each other)
- power-law (many nodes with few links, few nodes with many links).

4.3 Weighted links

Like in neural networks, links in the networks we wish to investigate are typically "weighted": they have different values for their strength. The connection strength of a (multi)link expresses the probability that a challenge will be successfully propagated across that link, i.e. that an agent in the output set would take up the challenge proposed by an agent in the input set. Link strength can be formalized most simply as the conditional probability P(B|A) that agent B would take on a challenge received from agent A.

This probability can depend on different properties of the communicative connection between agents, including: *trust* (the degree to which the receiving agent takes challenges coming via the link from the sending agent seriously), *power* (the degree to which the sending

[44] M.E.J. Newman, 'The Structure and Function of Complex Networks', *SIAM Review*, 2003, 167–256; S. N. Dorogovtsev and J. F. F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW* (Oxford University Press, USA, 2003); S. H Strogatz, 'Exploring Complex Networks', *Nature*, 410 (2001), 268–276.

agent can compel the sending agent to act), *friendship* (the degree to which the receiving agent is eager to do something for the sending agent independently of any power relationships), *friction* (the effort or energy that it costs to transmit the challenge across the link) and *bandwidth* (the detail and clarity with which information can be transmitted along the link). Link strength therefore may represent any and all of these properties. For example, in our Talking Nets model[45], it represents trust.

A simple implementation of the probabilistic interpretation can be found in "particle swarms"[46], a discrete model of spreading activation, in which the degree of activation of a node corresponds to the number of "particles" present in that node. The particles spread out of the node by randomly choosing one of the outgoing links, with a probability proportional to the link strength. Thus, strong links will typically propagate more particles than weaker links. If we interpret a particle crossing a link as a challenge that is successfully propagated across that link, then the link strength effectively denotes the probability of this happening.

However, note that the success of the crossing (i.e. whether the receiving agent effectively addresses the challenge) depends first of all on that agent itself, that is, its selection criteria for relevance, and its capabilities for processing the challenge. The link strength merely represents the social-communicative aspect of the process, namely the relationship between sending and receiving agent determining in how far the receiving agent is likely to pay attention to the other's challenges. If the two agents have a very weak relationship (like in two agents who don't know each other but who both have read the same Wikipedia page), the probability of one picking up a challenge left by the other is a priori very low, and will be determined almost exclusively by the properties of the challenge and of the agents rather than by their relationship. However, this low probability may be compensated by the fact that the multilink connects a very large number of agents, so that at least one agent in the output set is likely to pick up the challenge.

4.4 Processing in the agents

We assume that a challenge is represented by a vector or list of variables, some of which are specified as numbers, others that remain unspecified (denoted as #), e.g. (3, 1, #, -4, #, #, 0). These can be interpreted as deviations from the ideal value, which is conventionally set at zero. Thus, the "ideal" vector has the form (0, 0, 0, …, 0). The # means that the agent does not know the value of the deviation for that variable. When an agent receives such a challenge, its task is to bring the vector as close as possible to the ideal vector, by changing some of the values. To do that, it needs specific skills. Different agents will typically have different abilities. E.g. one agent may be able to reduce the value of the first element with 10 or less units, while another may be able to fill in a numerical value for a # in the second place.

The simplest way to model the activity of an agent is to represent its abilities as a look-up table or a set of production rules. This is a fixed system of transformations that take a particular type of input vector and transform it into a particular type of output vector. For example one of the rules in the agent might have the form: if the 3rd variable has a value ≤ 4, then set its value to 0; if its value > 4, then set its value to the old value minus 4.

This model can be easily extended into a sequential "search" algorithm. When an agent receives a challenge vector, it can apply all its existing rules in some (specified or random) order, with each rule producing a change in the vector. Each application of a rule can be interpreted as a "trial", an attempt to come closer to a resolution for the challenge. The

[45] Van Overwalle and Heylighen, 606.

[46] M.A. Rodriguez, 'Social Decision Making with Multi-relational Networks and Grammar-based Particle Swarms', in *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference On*, 2007, p. 39–39.

rules may be applied iteratively, so that a particular processing step is repeated until the provisional result satisfies some stopping condition. This stopping condition could represent either a (for the agent) satisfactory resolution of the challenge, or an exhaustion of the agent's capability to make further progress.

An example of such an implementation can be found in the classifier system model[47]. Note that this actually does not include separate agents, only separate rules for processing the "messages" = challenges. But a single rule ("classifier") can be interpreted as an elementary agent with just a single skill. The stigmergic coordination of rules in the classifier system therefore can be seen as a form of distributed intelligence or "society of mind" in which elementary agents collaborate towards a collective solution.

A more sophisticated agent model would be a neural network, where the initial challenge vector is presented as an activation vector to the input layer of the network, this activation then spreads through the network, and the modified challenge vector is read off from the final activation of the output layer. This allows us to use the different techniques from connectionist modeling, such as letting the weights of the neural network adapt to their usages, e.g. via Hebbian or delta learning. It is also allows us to give the agent a working memory (i.e. an internal state that affects future processing), by leaving some activation after the challenge has been processed. This remaining activation may then direct or "prime" the processing of a subsequent challenge, thus keeping a direct trace of previous processing (as distinguished from the indirect trace left in long-term memory as a change in the link weights).

An advantage of the network representation for the agent's cognitive architecture is that it makes the overall model *recursive*: a *network of networks*. This means that the nodes of the global network can be interpreted as local networks, which process information or tackle challenges using essentially the same distributed dynamics of challenge propagation. An example of such a model can be found in the Talking Nets simulation[48], which considers a "social" network of agents that are each modeled as "neural" networks. A similar recursion can be found in the "society of mind" model: here an agent can be interpreted as a stigmergically coordinated system of more primitive agents, while the overall distributed intelligence can be interpreted as a stigmergically coordinated system of more complex agents.

If we can extend the distributed architecture downwards to the internal structure of the agents, we can also extend it upwards, and consider agent-like subsystems of the global network. These consist of tightly linked agents that collaboratively process challenges in such a well-coordinated, synergetic manner that appear as if they formed a single, collective agent (or what is called an "actor" in political science). This may be a good model for organizations, such as firms, governments, NGOs, etc. On the other hand, considering several levels of super-agents and sub-agents is likely to make the overall model overly complex, especially if these are all represented as connectionist networks. A possible compromise may be to represent agents as networks only if the functioning of these agents is characterized by an adaptive intelligence. That means that basic hardware or software agents could be represented by simple production rules, while human agents might have to be represented by neural networks.

4.5 Learning in the distributed network

The similarity between the distributed hypernetwork of agents connected by multilinks and a neural network suggests that the distributed network should be able to learn by differentially

[47] Holland and others.
[48] Van Overwalle and Heylighen, 606.

strengthening or weakening its links. There are two fundamental learning mechanisms in neural networks: Hebbian learning and delta (or reinforcement) learning.

Hebbian is the simplest mechanism: it increases the weight of a link each time the nodes that the link connects have been co-activated (or activated in short succession). Since there is no equivalent in the challenge propagation paradigm for pure activation, it is not obvious how this could be implemented: two linked agents may be active at the same time, but while addressing different, independent challenges. In this case, there does not seem to be a good reason to strengthen the link, as the activity of the one does not anticipate or otherwise entail the activity of the other. The case for link strengthening becomes stronger when they are both active working on the same challenge, but since challenges typically change while they propagate, it will be difficult to find two challenges that are exactly the same. This difficulty may perhaps be circumvented by defining a similarity measure on challenges (with vectors, this is typically the cosine of the angle between the two vectors), and strengthening the link in proportion to the similarity. However, the rationale for doing this remains weak.

A better argument could be made for delta learning. This is a form of reinforcement learning[49] in which a link is rewarded if it brings the challenge vector closer to a resolution, and punished if it pushes it further away from the resolution. If a resolution would be modeled as a vector with all components equal to 0, then it should be easy to calculate in how far the transmission across a link has improved the state, by computing the length (distance from zero) of the vector before and after the transmission event. The link can then be strengthened by an amount proportional to the decrease in length (or weakened proportionally to the increase in length). 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.

Note that with this interpretation, the learning actually happens inside the agent, not in the link. But the practical result is that the link strength increases, and this seems the simplest way to model the process. Moreover, a self-organizing network may actually adapt its links to the way they are used, e.g. by maintaining "traces" of successful use similar to the pheromones that ants leave on successfully exploited paths[50]. This was the idea behind the original model for a global brain rooted in a "learning web" where commonly used links between web pages would be strengthened[51].

The network does not even need any sophisticated learning abilities to adapt to its usage. Links will be created or reinforced by such mundane activities like adding someone's phone and address to your list of contacts, bookmarking a site, joining a community or organization (and thus getting easier access to its members), or putting up signs so that people can more easily find a room. All these activities change the environment (medium) 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

[49] Florentin Woergoetter and Bernd Porr, 'Reinforcement Learning', *Scholarpedia*, 3 (2008), 1448 <doi:10.4249/scholarpedia.1448>.
[50] Heylighen, 253–280.
[51] Johan Bollen and Francis Heylighen, 'A System to Restructure Hypertext Networks into Valid User Models', *New Review of Hypermedia and Multimedia*, 4 (1998), 189–214 <doi:10.1080/13614569808914701>; Heylighen and Bollen, p. 439.

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…

What is as yet unclear with this interpretation is how strong the reinforcement will be for a given improvement. In neural networks, the relative size of this is determined by the "learning rate", a constant that is multiplied with the "delta", i.e. the difference between target outcome and actual outcome, to produce the change in the link strength.

5. Measuring the intelligence of the network

We now appear to have gathered the necessary ingredients to model a network of agents exhibiting distributed intelligence. But to determine whether the model is successful, we need a measure of intelligence, so that we can ascertain how efficient a particular network is in resolving the challenges that are posed to it. The first problem is in generating challenges for the network to resolve. The simplest approach seems to be to generate random challenge vectors, and to "submit" these challenges (probably in several copies) to random agents in the network. These agents can then select one or more of the submitted challenges to deal with, and ignore the others. After an agent deals with a challenge, the challenge will typically be relaxed somewhat (i.e. moved closer to the nil vector of pure relaxation). The agent then passes on the challenge to one or more further agents, via one or more of its outgoing multilinks, who may relax it further, then propagate it further, and so on.

The measure could calculate the average degree of "relaxation" of all the challenges submitted to the network, i.e. the reduction in their distance to the nil vector. This average relaxation would vary from moment to moment. Assuming that relaxation and propagation occur in discrete time steps, each "tick of the clock" would instantaneously change the state of all relaxation vectors. If the network is intelligent, the average distance to the nil vector should decrease with each time step. Note that given the inherent stochasticity of the agents' decisions, there is no guarantee that every single action will decrease this distance. Locally, some of the agents' actions may be erroneous (increasing the length of a particular challenge vector), or ineffectual (no change in the length). However, assuming that agents have an in-built preference for relaxation, and that many agents are dealing with many challenges in parallel, it seems likely that the average for all these challenges would decrease monotonically (if it does not over one time step, it almost surely will over many time steps). Also note that we cannot expect all challenges to be completely relaxed in the end. It is more realistic to assume that some (aspects of) challenges will remain irreducible. These unsolved challenges may continue to propagate as long as there are agents willing to try tackling them.

The intelligence of the network could then be equated with the speed of decrease in the average length for all challenge vectors present in the network: the more quickly the network can "relax" its various challenges, the smarter it is. This now allows us to compare many different variations of the model to see which ones exhibit most intelligence. Some obvious variations to investigate are:

- different criteria used by the agents to select which challenges to address
- different cognitive architectures[52] (e.g. production rules or neural networks) for the agents
- different hypergraph topologies (e.g. stigmergy vs. binary graphs)
- different network topologies (e.g. hierarchical, clustering, power-law, etc.)
- different rules for propagating the challenges across the links
- different learning rules for changing the link weights in the network

[52] C. Gershenson, 'Cognitive Paradigms: Which One Is the Best', *Cognitive Systems Research*, 5 (2004), 135–156.

Initially, we can build a simulation using a relatively straightforward model, e.g. a simple network with rule-based agents sitting in each node that tackle every challenge they receive. If the agent can fully resolve the challenge, the process stops there. Otherwise, the agent propagates the (partially relaxed) challenge to one or more of its outgoing connections, with a probability proportionately to the weight of their links. This simple model would already give us an idea of the overall challenge propagation dynamics.

It moreover would allow some simple variations in parameter values, e.g. by changing the number of outgoing copies of an unresolved challenge. Increasing the number of copies (e.g. propagating it to 5 rather than 2 of the outgoing connections) would at first sight increase the intelligence of the network, as more agents would be available to deal with a challenge, and therefore the probability would be higher that one of them would fully resolve the challenge. On the other hand, that means that agents would receive more challenges to deal with. Assuming that they can only tackle one challenge at a time, this may create a backlog of unresolved challenges ("information overload"), which would slow down the whole process. Thus, we can expect a trade-off between wider propagation increasing distributed intelligence and greater congestion decreasing intelligence. In a second stage, this problem may be reduced by a more intelligent "routing" of challenges to those agents that have the smallest backlog, as investigated in Gershenson's self-organizing bureaucracy model[53]. This routing can be controlled both via the link weights (e.g. by reducing the weight of links to chronically overloaded agents, cf. [54]) and via the agent selection criteria (e.g. by making the overloaded agents more selective in accepting new challenges).

This example is just meant to show how the challenge propagation paradigm can be explored starting from a relatively simple formalism and simulation, and then gradually elaborated into a more complex and realistic model of a "global brain" type of distributed intelligence. As the simulation becomes more realistic, its results should become more relevant for the real world, i.e. to understand how changes might be made to the organization of the Internet and its users so as to increase the distributed intelligence of the network.

6. Synthesis: components of the formal model

Let us summarize the model by mathematically defining all the components that together form our model of a distributed intelligence:

Challenges are vectors $\mathbf{c} = (c_1, c_2, \dots c_n) \in \mathbb{R}^n$. The variables c_i represent valences along different dimensions: $c_i > 0$ represents an opportunity for gain in aspect i of the present situation, $c_i < 0$ a danger for loss in aspect i. Resolving the challenge means collecting all the gains, while avoiding or suppressing all the losses.

A fully *resolved challenge* is represented by the nil vector: $\mathbf{0} = (0, 0, 0, \dots 0)$. Partial resolution, which we may call *relaxation*, means collecting some gains and/or suppressing some losses. The *degree of relaxation* is the opposite of the *intensity of the challenge.* This intensity can be measured by the length of the challenge vector:

[53] C. Gershenson, 'Towards Self-organizing Bureaucracies', *International Journal of Public Information Systems*, 1 (2008), 1–24.

[54] I. Kassabalidis and others, 'Swarm Intelligence for Routing in Communication Networks', in *Global Telecommunications Conference, 2001. GLOBECOM'01. IEEE*, 2001, VI, 3613–3617.

$$\| \mathbf{c} \| = \sqrt{\sum_{i=1}^{n} c_i^2} \geq 0,$$ with $\|\mathbf{0}\| = 0$ representing minimal intensity = maximal relaxation.

Agents are discrete elements a_i of the agent set: $A = \{a_1, \ldots, a_m\}$

Note that setting the nil vector as a representation of full resolution implies that all agents experience resolution in the same way, i.e. they all have the same needs or values. In a more realistic model[55], each agent a can be assumed to have its own need vector $\mathbf{n}_a \in \mathbb{R}^n$, which represents its ideal state of resolution. In that case, the intensity of the challenge can be measured by the length of the difference between the challenge and the need vector: $\| \mathbf{c} - \mathbf{n}_a \|$

An agent a performs a mapping from the initial, received challenge $\mathbf{c}$ to the processed challenge:

$$\mathbf{f}_a : \mathbb{R}^n \to \mathbb{R}^n\colon \mathbf{c} \to \mathbf{f}_a(\mathbf{c}) = (f_a(c_1), \ldots, f_a(c_m))$$

This mapping can be realized by a set of production rules: $c_i \to f_a(c_i)$, a Markov transition matrix, a neural network, a heuristic search algorithm through the challenge space, or some other information processing model. Processing normally relaxes the challenge: $\|\mathbf{f}_a(\mathbf{c})\| < \|\mathbf{c}\|$. The agent is rewarded for this processing by a "fitness gain/avoidance of fitness loss" equal to $\|\mathbf{c}\| - \|\mathbf{f}_a(\mathbf{c})\| > 0$. Agents that collect a lot of fitness points may see their influence in the network increase.

When agents are simultaneously confronted with different challenges $\mathbf{c}^a$, $\mathbf{c}^b$, … $\mathbf{c}^z$, they apply a *priority function* to estimate the importance of those challenges:

$$p\colon \mathbb{R}^n \to \mathbb{R}\colon (c^k_1, c^k_2, \ldots c^k_n) \to p(c^k_1, c^k_2, \ldots c^k_n)$$

The priority will normally be proportional to the intensity: $p(\mathbf{c}^k) \sim \|\mathbf{c}^k\|$ (the farther the challenge from the nil challenge, the higher its priority), but is likely to also depend on other features, like "difficulty", "clarity" or "surprise", that still need to be formalized.

Agents then select either the challenge $\mathbf{c}^k$ for which $p(\mathbf{c}^k)$ is maximal, i.e. $p(\mathbf{c}^k) \geq p(\mathbf{c}^m)$ for all $m \neq k$ (optimizing), or all challenges $\mathbf{c}^k$ for which $p(\mathbf{c}^k)$ is satisfactory, i.e. $p(\mathbf{c}^k) \geq s$, where s is a constant designating the minimum priority for the agent to address a challenge (satisficing). They apply their mapping function to process the selected challenge(s) into new challenge (s), while ignoring the other challenges $\mathbf{c}^m$ until the next time step.

Normally, we can assume that an agent is always confronted with several challenges, either directly by the "challenge producing engine" (playing the role of Nature, or the hand of God), or indirectly via propagation along its incoming links, as we will formalize below. Given that an agent is programmed to tackle the challenges with the highest priority p, we can see an agent as a "challenge-seeking" automaton, motivated to collect a maximum of fitness within the constraints of its situation (available challenges $\mathbf{c}$) and abilities (processing function $\mathbf{f}_a$).

[55] Eva Busseniers, *Hierarchical Organization Versus Self-organization*, GBI Working Paper, 2012

Agents A constitute the nodes of a *hypergraph* $H = (A, L)$ with L representing the multilinks between agents: $L = \{l_1, \ldots l_n\}$, $l_i = (l^{in}_i, l^{out}_i)$. The input and output sets defining the multilink are simply subsets of the set of agents:

$$l^{in}_i = \{a_1, \ldots, a_k\} \subseteq A, \quad l^{out}_i = \{a'_1, \ldots, a'_m\} \subseteq A$$

Two agents $a_j, a_k \in A$ are *connected* if there exists a multilink $l = (l^{in}, l^{out}) \in L$, such that:

$$a_j \in l^{in}, a_k \in l^{out}$$

Links can have a *weight* varying between 0 and 1: $w(l) \in [0, 1]$. This typically represents the probability that a challenge issued by one of the agents in the input set would reach one (any?) of the agents in the output set. The probability interpretation implies the additional normalization condition that the sum of the probabilities for outgoing links for a given agent a cannot be larger than one:

$$\sum_{l \in L^{out}(a)} w(l) \leq 1, \text{ with } L^{out}(a) = \{l \in L \mid a \in l^{in}\}$$

Challenge *propagation* occurs when the processed challenge $\mathbf{f}_a(\mathbf{c})$ is effectively accepted by at least one other agent a' to which a is connected via a link l. The number of agents that accept the challenge will depend on these agents' priority function and satisfaction criterion. This agent will process the challenge according to its own processing function $\mathbf{f}_{a'}$, and potentially propagate it further via its outgoing links to a further "layer" of agents.

Learning may occur by increasing the strength of a link proportionately to the success of a propagation along that link, which is measured as the reduction in the degree of relaxation:

$$w(l) \rightarrow w'(l) = w(l) + r.(\|\mathbf{c}\| - \|\mathbf{f}_a(\mathbf{c})\|)$$

Here, $r \in [0,1]$ is the learning rate, $\mathbf{c}$ the initial challenge transmitted via the link l, and $\mathbf{f}_a(\mathbf{c})$ the best of the challenges processed by one of the agents that received $\mathbf{c}$ via l. Note that the weight decreases if $\|\mathbf{f}_a(\mathbf{c})\| < \|\mathbf{c}\|$, i.e. if the challenge has actually become more challenging (instead of more relaxed).

The *intelligence* of the distributed network can be measured as the average relaxation rate over the challenges (which is the same as the average fitness gain made by the agents):

$$I = \frac{1}{n}\sum_{i=1}^{n}(\|\mathbf{c}_i\| - \|\mathbf{f}_a(\mathbf{c}_i)\|)$$

Since it is impossible to calculate $\mathbf{f}_a(\mathbf{c})$ for all challenges $\mathbf{c} \in \mathbb{R}^n$, I will typically be calculated on the basis of a finite (but large) sample of randomly generated challenges: $\{\mathbf{c}_1, \mathbf{c}_2, \ldots \mathbf{c}_n\}$. Rather than a homogeneous or Gaussian distribution of challenges, a more realistic random generator should produce a power law distribution of challenges around 0, i.e. many small deviations, few large ones:

$$P(c_i) \sim |c_i|^{-\alpha}, \text{ with } 0 < \alpha < 3$$

The same reasoning seems applicable for generating random need vectors for the different agents:

$$P(n_i) \sim |n_i|^{-\beta}, \text{ with } 0 < \beta < 3$$

7. Conclusion

The challenge propagation framework 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 a conceptual analysis of the main components of the framework, we have made a first attempt at formally defining these components and their relationships. This not only provides 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.

While the development of such a model will still require a lot of research, the present analysis makes the problem much more tractable, by subdividing it into separate elements (such as challenge, agent and multilink), structures (such as topology of the network, or organization of the agents), processes (such as challenge selection, challenge processing, and challenge propagation), and measures (such as degree of relaxation, priority of a challenge, and overall intelligence). Each of these components can be investigated and modeled relatively independently of the others, thus making the overall design modular. This allows a clear division of labor, so that different researchers can be responsible e.g. for modeling the processing function of the agents, modeling the network structure, modeling the learning processes, or modeling the selection of challenges—each working on its own component without interfering with the work of the others.

8. References

Adar, E., and L. A Adamic, 'Tracking Information Epidemics in Blogspace', in *Web Intelligence, 2005. Proceedings. The 2005 IEEE/WIC/ACM International Conference On*, 2005, pp. 207–214

Allen, D., *Getting Things Done: The Art of Stress-free Productivity* (Penguin Group USA, 2001)

Baars, B.J., 'The Global Workspace Theory of Consciousness', *The Blackwell Companion to Consciousness*, 2007, 236–246

Bechtel, W., and A. Abrahamsen, *Connectionism and the Mind: An Introduction to Parallel Processing in Networks.* (Basil Blackwell, 1991)

Ben-Jacob, E., I. Becker, Y. Shapira, and H. Levine, 'Bacterial Linguistic Communication and Social Intelligence', *Trends in Microbiology*, 12 (2004), 366–372

Bollen, Johan, and Francis Heylighen, 'A System to Restructure Hypertext Networks into Valid User Models', *New Review of Hypermedia and Multimedia*, 4 (1998), 189–214 <doi:10.1080/13614569808914701>

Busseniers, Eva, *Hierarchical Organization Versus Self-organization*, GBI Working Paper, 2012

Colombetti, G., 'Appraising Valence', *Journal of Consciousness Studies*, 8 (2005), 103–126 <http://polorovereto.unitn.it/%7Ecolombetti/docs/GC_AppraisingValence05.pdf>

Dehaene, S., M. Kerszberg, and J.P. Changeux, 'A Neuronal Model of a Global Workspace in Effortful Cognitive Tasks', *Proceedings of the National Academy of Sciences*, 95 (1998), 14529

Dorogovtsev, S. N., and J. F. F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW* (Oxford University Press, USA, 2003) <http://books.google.com/books?id=wGxLHYs6jNwC&printsec=frontcover&dq=Evolution+of+networks:+From+biological+nets+to+the+Internet+and+WWW&hl=en&src=bmrr&sa=X&ei=4zTWT977KoX80QXExLn1Aw&ved=0CDMQ6AEwAA>

Faller, Larry D., 'Relaxation Phenomenon (physics and Chemistry)', *Britannica Online Encyclopedia*, 2012 <http://www.britannica.com/EBchecked/topic/496927/relaxation-phenomenon> [accessed 16 January 2012]

Fredrickson, B. L, 'The Broaden-and-build Theory of Positive Emotions.', *Philosophical Transactions of the Royal Society B: Biological Sciences*, 359 (2004), 1367

Gabora, L. M., 'Meme and Variations: A Computational Model of Cultural Evolution', *D. Stein (Ed)*, 1993

Gallo, G., G. Longo, S. Pallottino, and S. Nguyen, 'Directed Hypergraphs and Applications', *Discrete Applied Mathematics*, 42 (1993), 177–201

Gallo, G., and M. G Scutella, 'Directed Hypergraphs as a Modelling Paradigm', *Decisions in Economics and Finance*, 21 (1998), 97–123

Gershenson, C., 'Cognitive Paradigms: Which One Is the Best', *Cognitive Systems Research*, 5 (2004), 135–156

---, 'Towards Self-organizing Bureaucracies', *International Journal of Public Information Systems*, 1 (2008), 1–24

Gigerenzer, G., and R. Selten, *Bounded Rationality: The Adaptive Toolbox* (the MIT Press, 2002)

Goertzel, B., *Creating Internet Intelligence: Wild Computing, Distributed Digital Consciousness, and the Emerging Global Brain* (Kluwer Academic/Plenum Publishers, 2002)

Hawkins, Jeff, and Sandra Blakeslee, *On Intelligence* (Henry Holt and Company, 2005)

Heath, C., and D. Heath, *Made to Stick: Why Some Ideas Survive and Others Die* (Random House, 2007)

Heylighen, Francis, *A Tale of Challenge, Adventure and Mystery: Towards an Agent-based Unification of Narrative and Scientific Models of Behavior*, ECCO Working Papers, 2012

---, 'A Cognitive-systemic Reconstruction of Maslow's Theory of Self-actualization', *Behavioral Science*, 37 (1992), 39–58 <doi:10.1002/bs.3830370105>

---, 'Accelerating Socio-technological Evolution: From Ephemeralization and Stigmergy to the Global Brain', in *Globalization as Evolutionary Process: Modeling Global Change*, Rethinking Globalizations, 10 (Routledge, 2008), p. 284

---, *Cognitive Systems*

---, 'Collective Intelligence and Its Implementation on the Web: Algorithms to Develop a Collective Mental Map', *Computational & Mathematical Organization Theory*, 5 (1999), 253–280 <doi:10.1023/A:1009690407292>

---, '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.* (Uchitel Publishing, 2011), pp. 274 – 289 <http://pcp.vub.ac.be/papers/GBconceptions.pdf>

---, 'Formulating the Problem of Problem-formulation', *Cybernetics and Systems*, 88 (1988), 949–957

---, 'Self-organization in Communicating Groups: The Emergence of Coordination, Shared References and Collective Intelligence', in *Complexity Perspectives on Language, Communication, and Society*, ed. by Angels Massip Bonet and Albert Bastardas (Springer, 2012) <http://pcp.vub.ac.be/Papers/Barcelona-LanguageSO.pdf>

---, 'Self-organization of Complex, Intelligent Systems: An Action Ontology for Transdisciplinary Integration', *Integral Review*, 2011 <http://pespmc1.vub.ac.be/Papers/ECCO-paradigm.pdf>

---, *Stigmergy as a Generic Mechanism for Coordination: Definition, Varieties and Aspects*, Ecco, 2011 <http://pespmc1.vub.ac.be/Papers/Stigmergy-WorkingPaper.pdf>

---, *The GBI Vision: Past, Present and Future Context of Global Brain Research*, GBI Working Papers, 2011

---, 'What Makes a Meme Successful? Selection Criteria for Cultural Evolution', in *Proc. 16th Int. Congress on Cybernetics* (Namur: Association Internat. de Cybernétique, 1998), p. 423–418

---, 'Why Is Open Access Development so Successful? Stigmergic Organization and the Economics of Information', in *Open Source Jahrbuch 2007*, ed. by B. Lutterbeck, M. Baerwolff & R. A. Gehring (Lehmanns Media, 2007), pp. 165–180 <http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf>

Heylighen, Francis, and Johan Bollen, 'Hebbian Algorithms for a Digital Library Recommendation System', in *International Conference on Parallel Processing* (Los Alamitos, CA, USA: IEEE Computer Society, 2002), p. 439 <doi:10.1109/ICPPW.2002.1039763>

Heylighen, Francis, and K. Chielens, ‘Evolution of Culture, Memetics’, *Encyclopedia of Complexity and Systems Science* (Springer, 2008), p. 10370 <doi:10.1007/978-0-387-30440-3_189>

Heylighen, Francis, and Clément Vidal, ‘Getting Things Done: The Science Behind Stress-Free Productivity’, *Long Range Planning*, 41 (2008), 585–605 <doi:10.1016/j.lrp.2008.09.004>

Holland, John H., Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard, *Induction: Processes of Inference, Learning, and Discovery* (MIT Press, 1989)

Kassabalidis, I., MA El-Sharkawi, RJ Marks, P. Arabshahi, AA Gray, and others, ‘Swarm Intelligence for Routing in Communication Networks’, in *Global Telecommunications Conference, 2001. GLOBECOM’01. IEEE*, 2001, VI, 3613–3617

Maslow, Abraham H., ‘Deficiency Motivation and Growth Motivation.’, 1955

---, *Motivation and Personality*, 2nd edn (New York: Harper & Row, 1970)

McLeod, P., K. Plunkett, and E. T Rolls, *Introduction to Connectionist Modelling of Cognitive Processes* (Oxford University Press Oxford, 1998)

Minsky, Marvin, *The Society of Mind* (Simon & Schuster, 1988)

Newell, A., and H. A Simon, *Human Problem Solving* (Prentice-Hall Englewood Cliffs, NJ, 1972)

Newman, M.E.J., ‘The Structure and Function of Complex Networks’, *SIAM Review*, 2003, 167–256

Nicolis, G., and I. Prigogine, *Self-organization in Nonequilibrium Systems: From Dissipative Structures to Order Through Fluctuations* (Wiley, New York, 1977)

Van Overwalle, Frank, and Francis Heylighen, ‘Talking Nets: A Multiagent Connectionist Approach to Communication and Trust Between Individuals.’, *Psychological Review*, 113 (2006), 606

Parunak, H. Van Dyke, ‘A Survey of Environments and Mechanisms for Human-human Stigmergy’, in *Environments for Multi-Agent Systems II* (Springer, 2006), pp. 163–186

Prigogine, I., and I. Stengers, *Order Out of Chaos: Man’s New Dialogue with Nature* (Bantam books, 1984)

Rodriguez, M.A., ‘Social Decision Making with Multi-relational Networks and Grammar-based Particle Swarms’, in *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference On*, 2007, p. 39–39

Strogatz, S. H, ‘Exploring Complex Networks’, *Nature*, 410 (2001), 268–276

Trewavas, A., ‘Aspects of Plant Intelligence’, *Annals of Botany*, 92 (2003), 1–20

Woergoetter, Florentin, and Bernd Porr, ‘Reinforcement Learning’, *Scholarpedia*, 3 (2008), 1448 <doi:10.4249/scholarpedia.1448>