LA-UR-98-489

CONF-980669--

TITLE: **SYMBIOTIC INTELLIGENCE: SELF-ORGANIZING KNOWLEDGE ON DISTRIBUTED NETWORKS, DRIVEN BY HUMAN INTERACTION**

AUTHOR(S):
Norman Johnson, T-3
Steen Rasmussen, EES-5
Cliff Joslyn, CIC-3
Luis Rocha, CIC-3
Steven Smith, CIC-8
Marianna Kantor, CIC-3

SUBMITTED TO: *Sixth International Conference on Artificial Life, University of California, Los Angeles, California, June 26–29, 1998*

Los Alamos

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

FORM NO. 836 R4
ST. NO. 2629 5/81

DISCLAIMER

DISCLAIMER

Symbiotic Intelligence: Self-Organizing Knowledge on Distributed Networks, Driven by Human Interaction

Norman JOHNSON[a], Steen RASMUSSEN[a,b], Cliff JOSLYN[a], Luis ROCHA[a], Steven SMITH[a] and Marianna KANTOR[a].

[a]Los Alamos National Laboratory
Los Alamos, New Mexico, 87545
U.S.A.

[b]Santa Fe Institute
1399 Hyde Park Road
Santa Fe, New Mexico, 87501
U.S.A.

Abstract

This work addresses how human societies, and other diverse and distributed systems, solve collective challenges that are not approachable from the level of the individual, and how the Internet will change the way societies and organizations view problem solving. We apply the ideas developed in self-organizing systems to understand self-organization in informational systems.

The simplest explanation as to why animals (for example, ants, wolves, and humans) are organized into societies is that these societies enhance the survival of the individuals which make up the populations. We as individuals contribute to, as well as adapt to, these societies because they make life easier in one way or another, even though we may not always understand the process, either individually or collectively. Despite the lack of understanding of the "how" of the process, our society during our existence as a species has changed significantly, from separate, small hunting tribes to a highly technological, globally integrated society.

We combine this understanding of societal dynamics with self-organization on the Internet (the Net). The unique capability of the Net is that it combines, in a common medium, our entire human-technological system in both breadth and depth: breadth in the integration of heterogeneous systems of machines, information and people; and depth in the detailed capturing of the entire complexity of human use and creation of information. When the full diversity of societal dynamics is combined with the accuracy of communication on the Net, a phase transition is argued to occur in problem solving capability.

Through conceptual examples, an experiment of collective decision making on the Net and a simulation showing the effect of noise and loss on collective decision making, we argue that the resulting symbiotic structure of humans and the Net will evolve as an alternative problem solving approach for groups, organizations and society. Self-organizing knowledge formation from this symbiotic intelligence exemplifies a new type of self-organizing system, one without dissipation and not constrained by limited resources.

1 Introduction

Our goal is to analyze and facilitate how people, in the process of accessing and using information on networks, create new knowledge without premeditation. We argue that the symbiotic combination of humans and smart networks will result in a previously unrealized capability of collective problem identification and solution. This capability will complement our existing centralized approach to problem solving. This symbiotic intelligence will greatly increase the success of organizations in achieving their goals, better utilizing their resources and preparing for the future. For the human society as a whole, this new resource will improve our quality of life and vitality.

Why are ants, wolves and humans organized into societies? The simplest answer is that these societies enhance the survival of individuals and, therefore, the populations in turn. Evolution has provided us with a method by which vital, complex problems of importance to the individuals that make up the societies are solved. Yet we, individually or collectively, do not understand the process. Despite the lack of understanding, our society has changed dramatically, from small hunting tribes to a highly technological, global society. (For a full discussion of system evolution see [Heylighen et al., 1993; Joslyn et al., 1995].)

One of the apparent paradoxes associated with human society is its diversity in both goals and means. Subgroups often tend to disagree significantly. Different individuals, different political parties, different companies, employers and employee, etc., create a complex network of contradicting means and goals. We shall argue that this divergence is not a liability or even an inconvenience, but is of vital importance and inherent survival value. Our society is a self-organizing system which involves many individuals and organizations at the "controls".

Information technology, from initially painting and writing, has been integral to extending our core social unit to global proportions. We shall argue that the growth of the Internet (the Net) is changing fundamental dynamics of human society's approach to the solution of complex problems. The timing of this new capability is of critical importance. We have created a highly volatile global society, but we have not, as yet, enabled the self-organizing dynamical system to give it stability and capability to manage itself.

In the remainder of the paper we first discuss a necessary addition to our traditional approach to problem identification and solving. From this perspective, we discuss the unique capabilities of the Net in integration and capturing the dynamics of human interactions. Finally we argue for the existence and usefulness of self-organizing knowledge formation on the Net.

The current study interacts with science in two ways. It is through the existing understanding of self-organizing material systems, particularly from the field of Artificial Life and computational biology, that the dynamical process of self-organizing behavior on the Net can be analyzed and developed. (We use "emergent" in the accepted usage of an observable property/structure of a system that cannot be detected or measured at the level of the structures or elements that make up the system. We use "self-organizing" in the restricted sense of a system that does not contain centralized control or global sense of itself.) The quantitative nature of interactions on the Net, as will be argued shortly, results in a rich medium to further the study of self-organizing systems. Secondly the new capability itself is expected to benefit scientific communities and endeavors by facilitating problem identification, particularly across discipline boundaries.

2 An Alternative Paradigm for Problem Solving

Most of our day-to-day technical developments and problem solving are achieved in a linear, premeditative approach:

- Understand the system in question,
- Represent the system behavior in an approximate model,
- Predict of the system behavior based on the model,
- Apply the model to optimize and control some aspects of the system to meet defined goals.
- Repeat the process for corrections or failures.

As long as we as a society of technocrats can "understand" the system, the above approach works well and we can optimize it for our needs. This approach is desirable because it has the potential to achieve optimal and rapid solutions.

Distributed systems (biological evolution, immune systems, human societies, etc.) take a very different approach to problem solving. When the system is too complex for analysis (as in global economics), an "understanding" of the whole is not possible. Without this understanding or when a centralized control is absent, the linear approach to optimization of a system is interrupted at the onset. Instead, a solution by distributed systems relies on a broad diversity of potential solutions at any time, combined with an emergent system dynamics. Then, given a change of state of the system - a new problem, one of the potentials is selected through the dynamics of the system. If a solution is not found within the current diversity, the system fails and "dies." If a possible solution is found, then the dynamics of the system will "optimize" the solution, but not by sacrificing the complete diversity of the system, typically by changing the relative dominance of subsystems. Thus, a solution is found without premeditation or conscious selection at some high level.

Herein lies the thrust of our idea. What we are suggesting will never replace the linear problem solving for less complex problems because of its ability to quickly optimize the solution space, typically requiring fewer resources. But for systems that defy timely understanding, due to the size or variety of information, only a distributed approach to problem solving can discover the solution space.

We strongly believe that in the future, emerging solutions on distributed networks will hold a respected place in problem solving and in science, once we accept the limits in our understanding and once the capability of emerging solutions is demonstrated. This way of solving problems has been around since the origins of life and now our technological capabilities allow us to further understand and improve this capability in a clear and precise manner by the use of the Net.

The view of human society as a collective organism is not new. George Dyson, in his recent book, "Darwin Among the Machines" surveys the works of great thinkers over the last 500 years who have touched on visions which involve evolution in the collective intelligence or awareness of humans and machines as symbol processors [dyson1997].

3 Why Now and Not Before

The Net has three significant capabilities beyond prior human-technological systems: (1) the ability to integrate quickly heterogeneous systems (the breadth of systems), (2) its ability to capture detailed signatures of the access and use of information (the depth of systems) and (3) its ability, with minimal loss, to relate and transmit information (the accuracy of communication).

1) The Net has the ability within one hyper-system to integrate:

a. Information storage, both in the form of simple data and complex text and images. This was earlier done in off-line libraries and a variety of data banks.

b. Communication. Communication was earlier done either by the relatively slow movement of people or documents or, in recent times, by telephone or other electronic technologies. However, complex documents, simple data and images can now be transported instantaneously and close to cost-free from anywhere to everywhere. Geographical barriers are virtually gone.

c. Traditional computing: the automated (simple) information processing of huge amounts of data.

d. Human processing. The human ability to analyze, understand and process limited, but highly complex information.

Until very recently (a), (b) and (c) were physically separated processes, all combined by human intervention (d). Now (a), (b) and (c) are integrated in a more standardized medium so (d) is no longer the slowest link and now can be active at a more efficient and useful level. Thus, the time scale for knowledge organization and creation is drastically shorter. This new integration has been overwhelming to humans, but tools are readily being developed in this infant hyper-structure to overcome this initial shortcoming.

2) The Net can capture the complexity (the depth) of how information is associated by retaining all references between data on the network. A simple example of how much of this relational information is currently lost is the use of scientific publications. While papers contain citations that connect the current paper with the information in other papers, the more immediate and detailed information about the numbers and types of readers of the papers can be only obtained in the past at great expense. Now with the advent of online publications, such information is explicitly available at effectively no cost. In general, the Net can capture all traces of the uses of information. As more and more of the human activities occur on the Net, these traces will capture the full

complexity of our interactions and activities that were formerly lost. While currently not used to any extent, these traces represent implicit knowledge of how we interact and how new knowledge is created.

3) Prior relational knowledge dissipates in traditional human-to-human communication, resulting in a rapid loss of information a bit removed from its creator (the children's game of passing a phrase around a circle is a telling example of the high noise-to-signal ratio of verbal communication). By contrast, information related on the Net, whether in web pages or emails, suffers no loss of information during transmission.

It is this loss of information that has lead to a maximum critical number of participants for effective problem solving in all human systems, from social functions to organizations. This inherent loss in traditional communication can prohibit problem definition, let alone the problem solution.

With the stronger presence of these three capabilities in human use and interactions on the Net, it can be argued that an informational "phase change" occurs: for the first time we can capture our creation, manipulation and rejection of knowledge, encompassing the full complexity of our existence - information which was earlier lost in the traditional person-to-person interactional noise or never available. With this phase change of information processing comes the possibility of an enhanced dynamic that, while an essential part of human organizations since our beginnings as social animals, has been limited in extent and responsiveness: that is, the growth of self-organizing knowledge in social systems.

4 Self-Organizing and Emergent Knowledge

Large problems are sometimes "solved" without our explicit intervention and may only be seen as problems in retrospect. A "solution" emerges and takes care of the problem without any single person identifying the problem or orchestrating the solutions.

A simple demonstration of a self-organizing solution is the formation of a path in a forest. The undefined goal is a path that is usable and of minimum length between two destinations and is adaptable to variations in local conditions. The self-organizing path evolves over time and includes contributions from a variety of people and animals with a broad diversity of goals (mode of travel, destinations, security, etc.) and changing forest conditions (weather, tree falls, dangerous locations, etc.). A comparable "designed" path would require a thoughtful process of premeditated selection of the goals to be optimized, an analysis of the forest layout and dynamics, a plan to meet the goals, and the implementation of the solution.

Self-organizing solutions have existed long before humans became capable of alternative approaches to problem solving. Even after our predominance as problem solvers arose, self-organizing solutions to our societal evolution continues to this day. The shift from an agricultural-based society to an industrial-based society occurred without any group of individuals being responsible. The society self-organized into another form.

Only in retrospect, it is clear that a number of key factors had to be in place to make such a shift possible. Only in retrospect, do we understand and model the process with the traditional process of knowledge creation.

Because prior studies of self-organizing systems have not been informational systems, some explanation of the borrowed nomenclature is necessary. We distinguish self-organized knowledge as being different in origin, not necessarily in kind, from traditional knowledge. We define traditional knowledge as being derived from the premeditative actions of people or predetermined processes of automated information processing to associate facts and draw conclusions. It represents a condensation of information and is localized in time and space - in books, speeches, memories. Self-organized knowledge, in contrast, originates from the self-organizing dynamics of hyper-systems, without the action of a central coordinating processor. Once self-organized knowledge is formed, it can either be observed as traditional knowledge or as emergent knowledge.

We now give two examples of how the Net might enhance the formation of self-organizing or collective knowledge in human interactions.

5 Demonstrations of Self-Organizing of Knowledge

We now introduce two systems which demonstrate collective knowledge development in different ways.

5.1 Adaptive Hypertext Experiment

A simple experiment was conducted by Bollen and Heylighen [1996a] of the Free University of Brussels under the Principia Cybernetica Project's (PCP) goal to to explore the "brain metaphor" [Gaines, 1994; Heylighen and Bolllen, 1996] to make hypertext webs more intelligent [Bollen and Heylighen, 1996b]. This metaphor led them to consider hypertext links like neural associations in the brain according to a Hebbian dynamics: "The strength of the links, like the connection strength of synapses, can change depending on the frequency of use of the link. This allows the network to 'learn' automatically from the way it is used" [Ibid], which illustrates the concept of emergent knowledge through human interaction.

The experiment was set up by first constructing a list of the 150 most common words in newspaper English. Initially when a user entered the system, a target word was displayed on a web page, followed by a list of 10 more randomly chosen words from the list (more words were available from the list without replacement at the user's request, to the point of potentially exhausting the list). The user is then asked to pick the word from the list that most closely is associated with the header word. Upon choosing a word, the order of the list is recalculated based on the frequency of selection according to a Hebbian rule, with weight added to the initial link, the reflexive link backwards, and the transitive link across two pairs of words. The user is then taken to a new page corresponding to the selected word, and the process is repeated. The researchers found that the lists stabilized to a fixed order after about 4000 selections in a site.

The resulting ordered lists determined a common semantics despite the heterogeneity of users. This simple task of ordering is easy for an individual but of little utility due to large individual variation in semantic differences in the community of web users. The network solution actively constructed useful collective knowledge representing a consensual semantics, but with minimal instruction and effort from the collective group of individuals. This example captures the essence of developing an self-organizing Knowledge system that combines the advantages of both human and computer networks to quickly solve a syntactically complex problem. From this example, one can imagine a host of previously challenging, if not intractable, problems that could be addressed once the methodology is developed.

One of the concerns one might have about automatic emergent knowledge generation, as pursued by this experiment, is the dependence on initial conditions. Given an initial sequence of 10 randomly chosen words, those words shown at the top of the list will tend to be reinforced by merely being shown first to users. This way, potentially, the successive chains of association might be biased into different "semantic attractors" depending on the initial list. To study this phenomenon in detail more extensive experiments than those so far conducted must be performed.

5.2 Simulation of Collective Decision Making

A complementary demonstration was done at Los Alamos [Johnson,1998] that supports some of the fundamental assumptions of the present argument and illustrates some of the features of a self-organizing system. We asked the question: what is the effect of noise or information loss on a collective decision involving many individuals. The general system that was examined was a maze (a connected, undirected graph) which has multiple minimal solutions between two nodes (one being the starting point and the other being the end or goal).

Many solutions (100s) to the maze were then found for a large group of independent "individuals" (there is no shared information between individuals as in the PCP example). All individuals use the same set of rules (the Learning Rules) that determine their movement through the maze, based only on local information, and how they modify their own path "preference" at each node. The restriction to using only local information means that they have no "global" sense of the maze and explore the maze until they just happen to reach the end node. The nodal path preference is a weighted, directed graph overlaying the maze and is retained for each individual for later use. Because random choices are made in the rules, a diversity of dominant paths through the maze and a diversity of lengths of paths ("performance") are created. Another set of rules (the Application Rules) then use, but do not modify, the nodal preferences to find the "optimal" path of each individual. Once the individual nodal preferences are found from the Learning Phase, these are combined to create a collective nodal preference, and the same set of the Application Rules are used to determine the collective solution.

For a demonstration maze of 35 nodes with 8 paths of a minimum path length of 9 (see Fig. 1), the average number of steps to "solve" the problem of 100 individuals is 34.3 (with a standard deviation of 24.5) in the Learning phase. The average performance of the individuals using the Application Rules is 12.8 with a standard deviation of 3.1. A simple average of the individual nodal preferences is used to create a collective nodal preference. Its application using the same Application Rules results in solution of 9-11 steps with more than 20 individuals, most often sampling the minimum path length. Figure 2 shows the change in path length as the numbers of individuals in the collective solution increases. The normalization of the collective path length is by the average individual path length. An average random walk solution of the example maze is 138 with a standard deviation of 101.

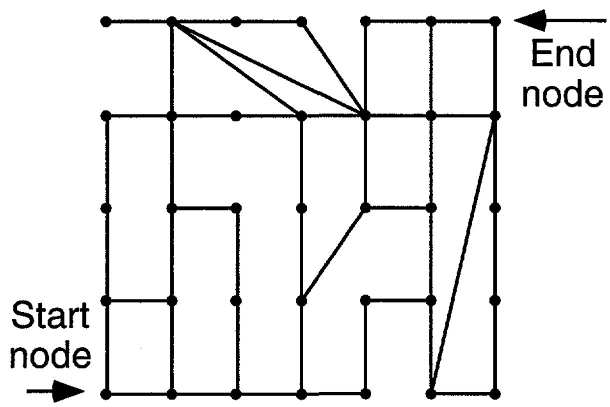

Figure 1. The "maze" used for the demonstration problem.

The significant improvement of the collective solution over the average individual solution illustrates, at least for this example, that information can be combined from uncoupled solutions to achieve a more optimal solution to a problem.

A few properties of the system illustrate some of the fundamental assumptions and arguments presented earlier. To support the observation that either noise or loss in communication is the cause of degradation of a collective decision in traditional ways of cooperation, the effects of noise and loss of the individual's contribution was examined. In general, it was found that the collective solution was remarkably robust. Degradation of the individual's contribution generally had no effect or just postponed the collective convergence to the minimal solution.

Some effects significantly interrupted the collective solution. The first was the random selection and use of the nodal preference of one of the contributing individuals, with a different individual selected at each node (see Fig. 2). The resulting average path length was about 45, independent of the number of individuals contributing to the solution. This is the worst case of leadership in a collective decision, but illustrates how the change of a dominant individual, in the absence of a collective decision, during a solution process can yield results much worse than that of an average individual. Likely the actual effect in a human interactions, as in a committee, is less severe, but probably a significant source of degraded performance. A second degradation of the collective solution was achieved by the random addition of noise to the contribution of an individual's nodal preference to the collective solution. This might be interpreted as miscommunication of the individual contribution to the whole. At moderate random addition, around half of the time and above, the collective solution does worse than the average individual performance. Hence, noise in the contribution of an individual, can degrade the collective solution below the performance of the individual.

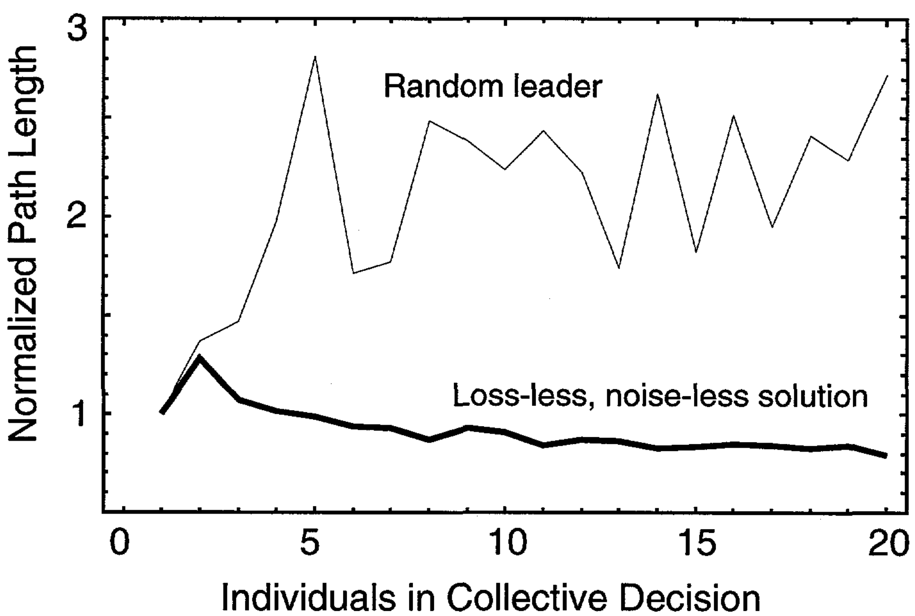

Figure 2. Plot of the path lengths of the collective solutions for with and without loss of information versus individuals contributing to the collective solution.

These results suggest that, as argued in Section 3, by removing sources of noise and loss in collective solutions, it might be possible to have larger numbers of individuals contribute to collective decision making than was previously possible.

Another observation was that diversity leads to robustness or improved contingency of the collective solution. For example, it was observed that the collective solution is degraded if only the "better" individuals (those with shorter path lengths) contribute to the collective solution, illustrating that even a diversity of performance is important to a collective solution. Another example is to apply the Learning phase to more than one goal (end node) for the group of individuals (i.e., each individual learns with one goal out of many) and then to find the performance of the average individual and collective in finding one of the specified goals. The collective far outperforms the average individual in this demonstration, averaging almost one-half of the individual solutions, as seen in Fig. 3. Another example was to change the goal (final node) after learning to different goal, measuring the robustness of the solution. Again, the collective decision was found to perform significantly better than the average performance of the individuals, averaging better than one-half of the individual solutions. These demonstrate that a collective solution from a diverse population is more flexible and performs better in changing goals than a single more narrowly-focused individual.

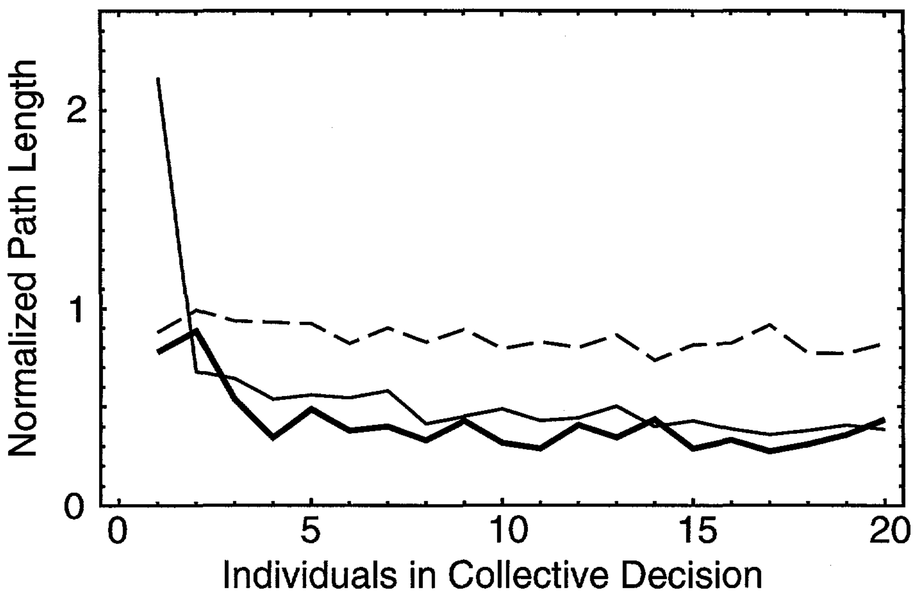

Figure 3. Plot of the path lengths of the collective solutions for three different node goals versus individuals contributing to the collective solution. The one curve that shows little improvement with increasing numbers of individuals is due to the individual and collective result being close to the optimum solution for the maze.

6 Conclusions

This paper presents preliminary arguments on the possible future of "problem solving" or collective decision making in our society and organizations. We have argued that the capability that underlies all life, the ability of self-organizing systems to "solve" essential problems, will take on new functionality as our society increasingly utilizes the Net for human interaction.

We suggest that self-organizing knowledge formation represents a self-organizing system that is unique: it is not dissipative, nor constrained by limited resources, as are prior examples of self-organizing systems. Knowledge has the unique property of non-conservation: the act of giving knowledge does not erase knowledge in the giver and new knowledge is not created at the expense of prior knowledge. Even knowledge found to be false is useful knowledge. This is in contrast to both biological and social self-organizing systems where success of one subsystem is generally at the expense of another.

These unique characteristics of self-organizing knowledge systems are further enabled by the unique capabilities of the Net: in addition to providing improved temporal and spatial system dynamics, the Net has noise-less and loss-less transmission of information, unlike our predominant verbal communication. The symbiotic intelligence of the combined human-Net system is believed to be able to operate at a level of functionality, both in numbers of individuals and the complexity of the capability, higher than previously possible.

To support this argument, we describe two demonstrations of collective intelligence. The PCP example of ordering word lists captures the creation of self-organizing knowledge by the interaction of humans processing complex semantic content, facilitated by the Net. This example illustrates the ease of solution to a problem that would be difficult using traditional approaches. The Los Alamos simulation of a simple system, with no interaction between individuals and no human participation, demonstrates the possibility of increased functionality in collective decisions with greater numbers of contributing individuals in the absence of noise and information loss in communication. While neither of these capture the full promise of the current argument, they capture two possible extremes of the potential of symbiotic intelligence.

References

[1] F Heylighen, C Joslyn and V Turchin 1993, "Super- and/or Meta-being(s)", http://pespmc1.vub.ac.be/SUP-META.html

[2] C Joslyn, V Turchin, and F Heylighen 1995, "Social Evolution ", http://pespmc1.vub.ac.be/SOCEVOL.html

[3] Dyson, G. [1997]. "Darwin among the Machines: The Evolution of Global Intelligence." Addison-Wesley Publishing,

[4] Bollen J, and F. Heylighen [1996a]. "Algorithms for the self-organization of distributed, multi-user neyworks. Possible application to the future World Wide Web." In: Cybernetics and Systems '96. R. Trappl (Ed.). Austrian Society For Cybernetics Press, pp. 911-916. (Also see, "The Adaptive Hypertext Experiment", http://pespmc1.vub.ac.be/ADHYPEXP.html)

[5] Bollen J, and F. Heylighen [1996b]. "Learning, "Brain-like" Webs". WWW document, http://pespmc1.vub.ac.be. (Also see, F Heylighen 1997, "Correspondence between Organism and Society", http://pespmc1.vub.ac.be/COMPTABL.html)

[6] Gaines, B.R. [1994]. "The collective stance in modelling expertise in individuals and organizations." International Journal of Expert Systems. Vol. 7, no. 1, pp. 22-51.

[7] Heylighen, F. and J. Bollen [1996]. "The World-Wide Web as a super-brain: from metaphor to model." In:Cybernetics and Systems '96. R. Trappl (Ed.). Austrian Society For Cybernetics Press, pp. 917-922.

[8] Johnson, N. L. [1998]. "Loss and Noise in Collective Decision Making, A Demonstration for the Symbiotic Intelligence Project," manuscript in preparation (contact the author at nlj@lanl.gov).