LA-UR-98-489

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

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*

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

Norman L. 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
[b]Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico, 87501

Abstract

Through conceptual examples and demonstrations, we argue that the symbiotic combination of the Internet and humans will result in a significant enhancement of the previously existing, self-organizing social structure of humans. The combination of the unique capabilities of intelligent, distributed information systems (the relatively loss-less transmission and capturing of detailed signatures) with the unique capabilities of humans (processing and analysis of complex, but limited, systems) will enable essential problem solving within our increasingly complex world. The capability may allow solutions that are not achievable directly by individuals, organizations or governments.

Introduction

The premise of our work is presented in this section. We acknowledge that the following ideas are still somewhat controversial within their own fields of relevance, but we are encouraged also by the growing integration of these ideas across many disciplines and are confident that the viewpoint presented here will be demonstrated and generally accepted.

The argument follows the path of (1) the evolution of human social behavior, (2) the effect of technology on social dynamics and structures, and (3) the relationship of system complexity and traditional problem solving within these social structures. These arguments lead us to our beginning point of the technology of the Internet (Net) and how it will change how humans solve problems.

We use "problem solving" in a broader context than the traditional usage of finding a solution to a problem by analysis. We include the ability of a dynamical system to "find" a new "solution" upon a change of state. While the usage can be problematic, no existing words/language seems suitable to cover both applications. The need of this inclusion will be apparent.

We start with the premise that we have evolved social structures, and the supporting dynamics, which enabled us to "solve" problems that threaten our existence (Joslyn, et al. 1995, Byron 1998). Unlike biological evolution, social change has the distinct advantage of enabling us to adapt within our own lifetime. Although possibly different in detail, social and biological evolution use the same dynamical processes and exhibit the same properties, inherent to self-organizing systems (see e.g., Babloyantz 1991, Forrest 1990 and the Artificial Life Proceedings I-V):

- "Solutions" arise as a selection by the system dynamics, driven by local processes, from a diversity of potential solutions. Selection does not typically reduce diversity, but only shifts the relative prevalence of the subsystems.
- These systems have the properties of distributed "control" (control from the bottom up), redundancy and persistent non-equilibrium.
- The global properties are: functionality greater than the individual subsystems, the capability to find solutions in the presence of conflicting needs, and scalability without loss of viability.

The view of human society as an adaptive, collective organism is not new. George Dyson (1997) in *Darwin Among the Machines* surveys the works of thinkers (e.g., Hobbes and Liebniz to Margulous) who have touched on this vision of society during the past five centuries. Despite the long history of interest in these ideas, it has only been in the last decades that there is now promise of a quantitative theory of social dynamics. This new foundation was driven by the dramatic success of the application of complex systems methods to biological problems as expressed, for example, in the Artificial Life movement. In the last two decades there has been a virtual explosion of interpretations or dynamical theories of social and economic systems (e.g., citations in Abraham 1994).

Evidence of our social evolution in action is easily seen in how we have adapted to the significant changes in technology, even though we are biologically unchanged for many millennia. The changes are most apparent in the dramatic increase in the maximum size of a social group as a result of technology advances in transportation, communication and knowledge storage. With each advance, the maximum size of a functioning social group has increased from initially tribes, to city-states, to nations, to regional coalitions, to finally global coalitions. These major societal shifts have occurred by processes similar to biological evolution without centralized planning, often with extreme diversity of capabilities and goals, and with solutions often far beyond the ability or understanding of any individual.

An central question at this juncture is "what is the role of individual or organizational problem solving within the context of self-organizing social dynamics?" Certainly many important societal shifts have resulted from the work or influence of a single individual, organization or government. Arguably these contributions may be necessary components to the overall dynamics, representing the actions of a mostly autonomous entity in a hierarchical self-organizing system.

But what is more important is that the capability of the individual, organization or government will falter, and possibly fail, if centralized problem solving is applied to a system that is not understandable. Without the understanding, there cannot be the analysis and prediction necessary for an effective and timely solution; there can only be trial and error. Humans are premiere problem solvers in systems with heterogeneous data of limited quantity, but we are overwhelmed by vast amounts of homogenous data. Obversely our computer processing counterparts are overwhelmed by complex data of any extent. Furthermore, we are limited in our ability to combine individual resources to solve problems of greater complexity, such as is observed in the limit on the maximum size of a useful committee.

If organizations or societies were to rely on just centralized control to solve problems, we would expect these efforts to fail as our society or the domain of our organizations becomes too complex. Social structures that take advantage of our inherent, self-organizing social dynamics will be best enabled to cope with our increasingly complex world (Abraham 1994). Indeed, we argue that this has happened in modern, overly centralized governments, such as the USSR, and is the reason that democracy and capitalism provide the most robust solutions in modern times (Slater and Bennis 1964 and 1990). There are also trends towards decentralized corporate management (Anderson and Arrow 1988, Youngblood 1996).

Herein lies our proposition and starting point. Self-organizing social dynamics has been an unappreciated positive force in our social development and has been significantly extended, at least in scope, by new technologies. At the same time, our culture and society are facing greater challenges due to the increasing complexity of our world, both in vastness and heterogeneity, possibly to the point of global disfunction. We argue that the Internet (Net) will change and enhance our social dynamics, to the point of becoming a significant resource for organizations and society as a whole. Once better understood, the consequence for management and governments will be an emphasis on encouraging diversity, increased access to information, and decentralized control.

The Unique Capabilities of the Net and its Effect on Social Dynamics

The Net has three significant, arguably unique, capabilities beyond prior human-technological systems:
(1) *The Net integrates the breadth of diverse systems*. It has the ability within one hyper-system to integrate (Schement and Lievrouw 1987):

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

b. *Communication*. Communication was done earlier 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. Thus, the time scale for knowledge organization and creation using traditional, non-self-organizing methods, is drastically shorter. The new integration has been overwhelming to humans, but tools are readily evolving in this infant hyper-structure to overcome the initial shortcoming [e.g., firefly.net, amazon.com, alexa.com].

(2) *The Net captures the depth of systems.* It can capture the complexity 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 in the use of scientific publications. While papers contain citations that connect a paper with other papers, the information about the numbers and types of readers of the papers could be only obtained in the past at great expense. With the advent of on-line publications, such information is explicitly available at effectively no cost. In general, the Net can capture all traces of the use of information. These traces represent implicit knowledge of how we interact and how new knowledge is created. As (1) above is better realized, these traces will capture the full complexity of our interactions.

(3) *The Net has accuracy of communication.* Traditional human-to-human communication results in a rapid loss of information a bit removed from its creator (the children's game of whispering a phrase around a circle is a telling example of the high noise-to-signal ratio of verbal communication). By contrast, information exchanged or related on the Net suffers minimal loss of information during transmission or linking, in the same way that the

content of a book is not altered when exchanged. We do note that we sacrifice bandwidth using current technologies because of the elimination of vocal, facial and gestural expressions. In this discussion, we do not include the misinterpretation that can still occur in understanding of exchanged information; this source of miscommunication occurs regardless of the mechanism of exchange.

With the stronger presence of these unique capabilities of the Net in human dynamics, we propose that minimally the creation, manipulation and rejection of knowledge can be captured for the first time, encompassing the full complexity of the cognition process in our society. More importantly, the processes of our social dynamics, which previously relied on slower, spatially concentrated, and noisy forms of communication, now has the potential to form a symbiotic relationship between humans and the Net, enabling our prior self-organizing capabilities to operate at a significantly enhanced functionality. In the next section we give two examples of demonstrations of how this symbiosis might be possible. Furthermore, in the same manner as to how society self-organized to solve problems of survival, the same processes on the Net will result in self-organization of knowledge. Because self-organizing knowledge arises from diverse contributions and can encompass knowledge greater than the contribution of any individual, there is the arguable potential of creating knowledge that will contribute to solutions that are not understandable within our current processes. In the next section, we will also give a suggestive example of this capability.

Self-Organizing Systems Demonstrations

We now present two studies that demonstrate collective knowledge development: the first demonstrating knowledge formation from humans interacting on a network and the second examining how many individual solutions can combine to solve a global problem in an idealized system without human involvement.

Self-Organization on Networks: 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 goal to explore the "brain metaphor" (Gaines 1994; Heylighen and Bollen 1996) to make hypertext webs more intelligent (Drexler 1991, 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. When a user initially enters the system, a target word is 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 between individuals. 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 a 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.

Simulation of Collective Decision Making

The second demonstration is not an example of self-organization on an existing network, but a demonstration (Johnson 1998) that supports some of the fundamental assumptions of the present argument and illustrates desirable features of a large and diverse self-organizing system. We want to answer the following question: "what is the effect of noise or information loss on a collective decision involving many individuals."

The system that was examined was a maze (a connected, undirected graph) which has one or more solutions (paths) between two nodes (one being the starting point and the other being the end or goal). Solutions to the maze were found for a large (100s) number of independent "individuals" (no information is shared between individuals as in the prior demonstration). All individuals initially use the same set of "Learning Rules" that (1) determine their movement through the maze as based only on local information, and (2) 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 set of "Nodal Path Preferences" is a weighted, directed graph overlaying the maze and is retained for each individual for later use. Basically the Learning Rules select a link that has not been tried and then sets the Path Preference of this choice to be larger than the other links at this node. After the Learning phase is completed, another set of rules, the "Application Rules," are used. These apply, but do not modify, the nodal preferences to find the "optimal" path of each individual. Basically the Application Rules select the preferred link at a node with minor additional logic to prevent infinite loops. Because random choices are made in the rules between equal preference, a diversity of preferred paths through the maze and a diversity of total lengths of paths ("performance") are created. Once the individual nodal preferences are found from the Learning Phase, these can be combined at each node in various ways to create a collective nodal preference, and then the same set of the Application Rules are used to determine the collective solution.

For a demonstration maze of 35 nodes with 14 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. There is no correlation observed between the performance in the Learning and Application phases: a slow learner is not necessarily a poor performer. For the reference simulation, a simple average of the individual nodal preferences is used to create a collective nodal preference. Its application using the identical Application Rules results in solution of 9-11 steps when more than 20 individuals are included, most often sampling one of the minimum path lengths. Figure 2 shows the change in path length as the numbers of individuals in the collective solution increases. Note the effect of randomness, even though the identical individuals contribute to each collective decision. The average random walk solution is 138 with a standard deviation of 101. We note that the primary source of variation at larger contributors is due to the multiple minimum paths. Had there been only one minimum path, the solution is much more stable.

Figure 1. The "maze" used for the demonstration problem. Two of the 14 paths of minimum length are highlighted.

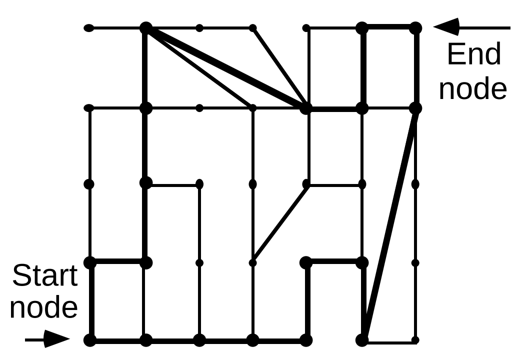

A few properties of the system illustrate some of the fundamental assumptions and arguments presented earlier. The significant improvement of the collective solution over the average individual solution (9 versus 12.8) illustrates that information can be combined from uncoupled individual solutions using only local information to achieve an optimal global solution to a problem. This emergent property of the collective system was generally observed on all mazes, even ones of higher complexity, with only difference being that different numbers of individuals are needed in the collective solutions to achieve the same performance.

In general, the collective solution was remarkably robust. Degradation of the individual's contribution, however implemented, generally had no effect or just postponed the collective convergence to the minimal solution. A few effects were found to significantly degrade the collective solution. One was the random selection and use of the nodal preference of one of the contributing individuals, with a different individual selected at each node. The resulting average path length was about 45 steps (3.5 normalized), independent of the number of individuals contributing to the solution. This illustrates how the change of a dominant individual during a solution process can yield results much worse than that of an average individual. A second degradation of the collective solution was achieved by the random addition of noise (the random replacement of a nodal preference by a small value) to the collective solution, in an attempt to model miscommunication of the individual contribution to the whole. At moderate random addition, around half of the time and greater, the collective solution does worse than the average individual performance. These results support the argument proposed in the prior section: many more individuals can contribute to a collective decision when sources of noise and loss are reduced.

Figure 2. Plot of normalized path lengths of the collective solutions versus the number of individuals contributing to the collective for two initial random seeds in the Application phase. The normalization is by the average individual path or about 12.8 steps.

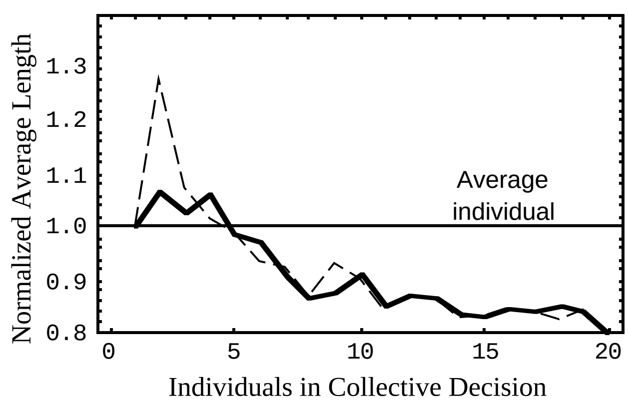

Another observation was that a collective solution from a diverse population is more flexible and performs better in changing goals than the average, more narrowly-focused individuals. For example, it was observed that the collective solution is degraded if only the "better" individuals (those with shorter path lengths in the Application Phase) 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 (i.e., each individual learns with one goal out of many) or to change the goal after learning with different goal, measuring the robustness of the solution. In both of these simulations, the collective decision performed significantly better with a normalized path of about 0.5.

There are obvious similarities between the processes we are describing here and what is being studied under the terms *Genetic Algorithms* and *Programming* (Koza 1994, Mitchell 1996). However, there are also some significant differences, perhaps the most important being that these agents do not evolve, but learn and create knowledge as they share information among themselves. The key to performance in these systems is diversity and not selection.

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 a dynamic process 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. 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 capability, higher than previously possible.

To support this argument, we have described two demonstrations of collective intelligence. The hypertext 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 demonstrates (1) the potential for more individuals to contribute to a collective solution, (2) the collective solution has better performance and is more robust than an average individual's solution, and (3) more complex problems can be solved with larger numbers of contributing individuals.

References

Abraham, R. 1994. *Chaos, Gaia, Eros*. San Francisco: Harper.

Anderson, P.H. & Arrow, K.A. 1988. *The Economy as an Evolving Complex System.* Reading, Mass.: Addison-Wesley Publishing.

Babloyantz, A. (Ed.) 1991. *Reviews and Commentary for Self-Organization, Emerging Properties, and Learning.* Portland, Or.: Book News, Inc.

Bollen, J. and Heylighen, F. 1996a. Algorithms for the self-organization of distributed, multi-user networks, Possible application to the future World Wide Web. *Cybernetics and Systems '96*. R. Trappl (Ed.). Austrian Society For Cybernetics Press :911-916.

Bollen J, & Heylighen, F. 1996b. *Learning, "Brain-like" Webs.* http://pespmc1.vub.ac.be.

Byron, M. 1998. Crisis-Driven Evolutionary Learning: Conceptual Foundations and Systematic Modeling. *World Congress on Sociology '98.* Montreal. Forthcoming.

Drexler, K.E. 1991. Hypertext Publishing and the Evolution of Knowledge. *Social Intelligence* 1:2-22.

Dyson, G. 1997. *Darwin among the Machines: The Evolution of Global Intelligence*. Reading, Mass.: Addison-Wesley Publishing.

Forrest, S. (Ed.) 1990. *Emergent Computation*. Cambridge, MA: MIT Press. (Special issue of Physica D)

Gaines, B.R. 1994. The Collective Stance in Modeling Expertise in Individuals and Organizations. *International Journal of Expert Systems* 7(1):22-51.

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

Johnson, N. L. 1998. *Effects of Complexity, Noise and Loss in Collective Decision Making.* Forthcoming.

Joslyn, C., Turchin V. , & Heylighen, F. 1995. *Social Evolution*, http://pespmc1.vub.ac.be/SOCEVOL.html.

Koza, J. 1994. *Genetic programming II*. Cambridge, MA: MIT Press.

Mitchell, M. 1996. *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press.

Schement, J.R., & Lievrouw, L.A. (Eds.) 1987. *Competing visions, complex realities: Social aspects of the information society*. Norwood, NJ: Ablex.

Slater, P. & W. Bennis 1964. Democracy is Inevitable. *Harvard Business Review*: March-April (reissued September–October 1990).

Youngblood, M. 1996. *Life at the Edge of Chaos*, New York: John Wiley & Sons.