Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map

Francis HEYLIGHEN*

Center "Leo Apostel", Free University of Brussels

Address: Krijgskundestraat 33, B-1160 Brussels, Belgium
E-mail: fheyligh@vub.ac.be
home page: http://pespmc1.vub.ac.be/HEYL.html

ABSTRACT.

Collective intelligence is defined as the ability of a group to solve more problems than its individual members. It is argued that the obstacles created by individual cognitive limits and the difficulty of coordination can be overcome by using a collective mental map (CMM). A CMM is defined as an external memory with shared read/write access, that represents problem states, actions and preferences for actions. It can be formalized as a weighted, directed graph. The creation of a network of pheromone trails by ant colonies points us to some basic mechanisms of CMM development: averaging of individual preferences, amplification of weak links by positive feedback, and integration of specialised subnetworks through division of labor. Similar mechanisms can be used to transform the World-Wide Web into a CMM, by supplementing it with weighted links. Two types of algorithms are explored: 1) the co-occurrence of links in web pages or user selections can be used to compute a matrix of link strengths, thus generalizing the technique of "collaborative filtering"; 2) learning web rules extract information from a user's sequential path through the web in order to change link strengths and create new links. The resulting weighted web can be used to facilitate problem-solving by suggesting related links to the user, or, more powerfully, by supporting a software agent that discovers relevant documents through spreading activation.

1. Introduction

With the growing interest in complex adaptive systems, artificial life, swarms and simulated societies, the concept of "collective intelligence" is coming more and more to the fore. The basic idea is that a group of individuals (e.g. people, insects, robots, or software agents) can be smart in a way that none of its members is. Complex, apparently intelligent behavior may emerge from the synergy created by simple interactions between individuals that follow simple rules.

---

* Research Associate FWO (Fund for Scientific Research-Flanders)

To be more accurate we can define intelligence as *the ability to solve problems.* A system is more intelligent than another system if in a given time interval it can solve more problems, or find better solutions to the same problems. A group can then be said to exhibit collective intelligence if it can find more or better solutions than the whole of all solutions that would be found by its members working individually.

1.1. Examples of collective intelligence

All organizations, whether they be firms, institutions or sporting teams, are created on the assumption that their members can do more together than they could do alone. Yet, most organizations have a hierarchical structure, with one individual at the top directing the activities of the other individuals at the levels below. Although no president, chief executive or general can oversee or control all the tasks performed by different individuals in a complex organization, one might still suspect that the intelligence of the organization is somehow merely a reflection or extension of the intelligence of its hierarchical head.

This is no longer the case in small, closely interacting groups such as soccer or football teams, where the "captain" rarely gives orders to the other team members. The movements and tactics that emerge during a soccer match are not controlled by a single individual, but result from complex sequences of interactions. Still, they are simple enough for an individual to comprehend, and since soccer players are intrinsically intelligent individuals, it may appear that the team is not really more intelligent than its members.

Things are very different in the world of social insects (Bonabeau et al. 1997; Bonabeau & Theraulaz 1994). The way that ants map out their environment, that bees decide which flower fields to exploit, or that termites build complex mounds, may create the impression that these are quite intelligent creatures. The opposite is true. Individual insects have extremely limited information processing capacities. Yet, the ant nest, bee hive or termite mound as a collective can cope with very complex situations.

What social insects lack in individual capabilities, they seem to make up by their sheer numbers. In that respect, an insect collective behaves like the self-organizing systems studied in physics and chemistry (Bonabeau et al. 1997): very large numbers of simple components interacting locally produce global organization and adaptation. In human society, such self-organization can be found in the "invisible hand" of the market mechanism. The market is very efficient in allocating the factors of production so as to create a balance between supply and demand (cf. Heylighen 1997). Centralized planning of the economy to ensure the same balanced distribution would be confronted with a "calculation problem" so complex that it would surpass the capacity of any information processing system. Yet, an efficient market requires its participating agents to follow only the most simple rules. Simulations have shown that even markets with "zero intelligence" traders manage to reach equilibrium quite quickly (Gode & Sunder 1993).

The examples we discussed show relatively low collective intelligence emerging from highly intelligent individual behavior (football teams) or high collective intelligence emerging from "dumb" individual behavior (insect societies and markets). The obvious question is whether high collective intelligence can also emerge from high individual intelligence. Achieving this is everything but obvious, though. The difficulty is perhaps best illustrated by the frustration most people experience with committees and meetings. Bring a number of very competent people together in a room in order to devise a plan of action, tackle a problem or reach a decision. Yet, the result you get is rarely much better than the result you would have got if the different participants had tackled the problem individually. Although committees are obviously important and useful, in practice it appears difficult for them to realize their full potential. Let us therefore consider some of the main impediments to the emergence of collective intelligence in human groups.

1.2. Obstacles to collective intelligence

First, however competent the participants, their individual intelligence is still limited, and this imposes a fundamental restriction on their ability to cooperate. Although an expert in his own field, Mr. Smith may be incapable to understand the approach proposed by Ms. Jones, whose expertise is different. Even if we assume that Mr. Smith would be able to grasp all the ramifications and details of Ms. Jones's proposal, he probably would still misunderstand what she is saying, simply because he interprets the words she uses in a different way than the one she intended. Both verbal and non-verbal communication are notoriously fuzzy, noisy and dependent on the context or frame of reference. Even if everyone would perfectly understand everyone else, many important suggestions during a meeting would never be followed up. In spite of note taking, no group is able to completely memorize all the issues that have been discussed.

Another recurrent problem is that people tend to play power games. Everybody would like to be recognized as the smartest or most important person in the group, and is therefore inclined to dismiss any opinion different from his or her own. Such power games often end up with the establishment of a "pecking order", where the one at the top can criticize everyone, while the one at the bottom can criticize no one. The result is that the people at the bottom are rarely ever paid attention to, however smart their suggestions. This constant competition to make one's voice heard is exacerbated by the fact that linguistic communication is *sequential*: in a meeting, only one person can speak at a time.

It seems that the problem might be tackled by splitting up the committee into small groups. Instead of a single speaker centrally directing the proceedings, the activities might now go on in parallel, thus allowing many more aspects to be discussed simultaneously. However, now a new problem arises: that of *coordination*. To tackle a problem collectively, the different subgroups must keep close contact. This implies a constant exchange of information so that the different groups would know what the others are doing, and can use each other's results. But this again creates a great information load, taxing both the communication channels and the individual cognitive systems that must process all this incoming information. Such load only becomes larger as the number of participants or groups increases.

For problems of information transmission, storage and processing, computer technologies may come to the rescue. This has led to the creation of the field of Computer-Supported Cooperative Work (CSCW) (see e.g. Smith 1994), which aims at the design of Groupware or "Group Decision Support Systems". CSCW systems can alleviate many of the problems we enumerated. By letting participants communicate anonymously via the system it can even tackle the problem of pecking order, so that all contributions get an even opportunity to be considered. However, CSCW systems are typically developed for small groups. They are not designed to support self-organizing collectives that involve thousands or millions of individuals.

But there is a technology which can connect those millions: the global computer network. Although communities on the Internet appear to self-organize more efficiently than communities that do not use computers, the network seems merely to have accelerated existing social processes. As yet, it does not provide any active support for collective intelligence. The present paper will investigate how such a support could be achieved, first by analysing the mechanisms through which collective intelligence emerges in other systems, then by discussing how available technologies can be extended to implement such mechanisms on the network.

2. Collective Problem-Solving

To better understand collective intelligence we must first analyse intelligence in general, that is, the ability to solve problems. A *problem* can be defined as a difference between the present situation, as perceived by some agent, and the situation desired by that agent. Problem-solving then means finding a sequence of actions that will transform the present state via a number of intermediate states into a goal state. Of course, there does not need to be a single, well-defined goal: the agent's "goal" might be simply to get into any situation that is more pleasant, interesting or amusing than the present one. The only requirement is that the agent can *distinguish* between subjectively "better" (preferred) and "worse" situations (Heylighen 1988, 1990).

To generalize this definition of a problem for a collective consisting of several agents it suffices to aggregate the desires of the different agents into a collective preference and their perceptions of the present situation into a collective perception. In economic terms, the aggregate desire becomes the market "demand" and the aggregate perception of the present situation becomes the "supply" (Heylighen, 1997). It must be noted, though, that what is preferable for an individual member is not necessarily what is preferable for a collective (Heylighen & Campbell, 1995): in general, a collective has emergent properties that cannot be reduced to mere sums of individual properties. (Therefore, the aggregation mechanism will need to have a non-linear component.) In section 3, we will discuss in more detail how such an aggregation mechanism might work.

On way to solve a problem is by trial-and-error in the real world: just try out some action and see whether it brings about the desired effect. Such an approach is obviously inefficient for all but the most trivial problems. Intelligence is characterised by the fact that this exploration of possible actions takes place mentally, so that actions can be selected or rejected "inside one's head", before executing them in reality. The more efficient this mental exploration, that is, the less trial-and-error needed to find the solution, the more intelligent the problem-solver.

2.1. Mental maps

The efficiency of mental problem-solving depends on the way the problem is *represented* inside the cognitive system (Heylighen 1988, 1990). Representations typically consist of the following components: a set of problem states, a set of possible actions, and a preference function or "fitness" criterion for selecting the most adequate actions. The fitness criterion, of course, will vary with the specific goals or preferences of the agent. Even for a given preference, though, there are many ways to decompose a problem into states and actions. Changing the way a problem is represented, by considering different distinctions between the different features of a problem situation, may make an unsolvable problem trivial, or the other way around (Heylighen 1988, 1990).

Actions can be represented as operators or transitions that map one state onto another one. A state that can be reached from another state by a single action can be seen as a neighbor of that state. Thus, the set of actions induces a topological structure on the set of states, transforming it into a problem *space*. The simplest model of such a space is a network, where the states correspond to the nodes of the network, and the actions to the edges or links that connect the nodes. The selection criterion, finally, can be represented by a preference function that attaches a particular weight to each link. This problem representation can be seen as the agent's *mental map* of its problem environment.

A mental map can be formalized as a weighted, directed graph $M = \{N, L, P\}$, where $N = \{n_1, n_2, ..., n_m\}$ is the set of nodes, $L \quad N \times N$ is the set of links, and P: $L \rightarrow [0, 1]$, is the preference function. A problem solution then is a connected path

$C = (c_1, ..., c_k)$ N such that c_1 is the initial state, c_k is a goal state, and for all i $\{1, \ldots, k\}$: (c_i, c_{i+1}) L.

To solve a problem, you need a general heuristic or search algorithm, that is, a method for selecting a sequence of actions that is likely to lead as quickly as possible to the goal. If we assume that the agent has only a local awareness of the mental map, that is, that the agent can only evaluate actions and states that are directly connected to the present state, then the most basic heuristic it can use is some form of "hill-climbing" with backtracking. This heuristic works as follows: from the present state choose the link with the highest weight that has not been tried out yet to reach a new state; if all links have already been tried, backtrack to a state visited earlier which still has an untried link; repeat this procedure until a goal state has been reached or until all available links have been exhausted. The efficiency of this method will obviously depend on how well the nodes, links and preference function reflect the actual possibilities and constraints in the environment.

The better the map, the more easily problems will be solved. Intelligent agents, then, are characterized by the quality of their mental maps, that is, by the knowledge and understanding they have of their environment, their own capacities for action, and their goals. Increasing problem-solving ability will generally require two complementary processes: 1) enlarging the map with additional states and actions, so that until now unimagined options become reachable; 2) improving the preference function, so that the increase in total options is counterbalanced by a greater selectivity in the options that need to be explored to solve a given problem.

2.2. Coordinating individual problem-solutions

Let us apply this conceptual framework to collective problem-solving. Imagine a group of individuals trying to solve a problem together. Each individual can explore his or her own mental map in order to come up with a sequence of actions that constitutes part of the solution. It would then seem sufficient to combine these partial solutions into an overall solution. Assuming that the individuals are similar (e.g. all human beings or all ants), and that they live in the same environment, we may expect their mental maps to be similar as well. However, mental maps are not objective reflections of the real world "out there": they are individual constructions, based on subjective preferences and experiences (cf. Heylighen 1999). Therefore, the maps will also be to an important degree different.

This diversity is healthy, since it means that different individuals may complement each others' weaknesses. Imagine that each individual would have exactly the same mental map. In that case, they would all find the same solutions in the same way, and little could be gained by a collective effort. (In the best case, the problem could be factorized into independent subproblems, which would then be divided among the participating individuals. This would merely speed up the problem-solving process, though; it would not produce any novel solutions).

Imagine now that each individual would have a different mental map. In that case, individuals would need to communicate not only the (partial) solutions they have found, but the relevant parts of their mental maps as well, since a solution only makes sense within a given problem representation. This requires a very powerful medium for information exchange, capable of transmitting a map of a complex problem domain. Moreover, it requires plenty of excess cognitive resources from the individuals who receive the transmissions, since they would need to parse and store dozens of mental maps in addition to their own. Since an individual's mental map reflects that individual's total knowledge, gathered during a lifetime of experience, it seems very unlikely that such excess processing and storage capacity would be available. If it were, this would mean that the individual has used only a fraction of his or her capacities for cognition, and this implies an in-

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.