

The World-Wide Web as a Super-Brain: from metaphor to model

Francis Heylighen & Johan Bollen*

Center "Leo Apostel", Free University of Brussels,
Pleinlaan 2, B-1050 Brussels
Belgium

email: fheylich@vnet3.vub.ac.be, jbollen@vnet3.vub.ac.be

Abstract

If society is viewed as a super-organism, communication networks play the role of its brain. This metaphor is developed into a model for the design of a more intelligent global network. The World-Wide Web, through its distributed hypermedia architecture, functions as an "associative memory", which may "learn" by the strengthening of frequently used links. Software agents, exploring the Web through spreading activation, function as problem-solving "thoughts". Users are integrated into this "super-brain" through direct man-machine interfaces and the reciprocal exchange of knowledge between individual and Web.

1 Introduction

It is a recurrent idea that the whole of humanity, the system formed by all people together with their channels of exchange, can be viewed as a single organism: the 'super-being' [Turchin, 1977] or 'metaman' [Stock, 1993]. When considering the remaining conflicts and misunderstandings, though, it becomes clear that the integration of individuals in human society is much less advanced than the integration of cells in a multicellular organism. Analysis of the evolutionary mechanisms underlying selfishness, competition and cooperation among individuals and groups moreover points to fundamental obstacles hindering further integration [Heylighen & Campbell, 1995].

Yet, there is at least one domain where integration seems to be moving full speed ahead: the development of ever more powerful communication media. In the *society as super-organism* metaphor, the communication channels play the role of nerves, transmitting signals between the different organs and muscles [Turchin, 1977]. In more advanced organisms, the nerves develop a complex mesh of interconnections, the brain, where sets of incoming signals are integrated and processed. After the advent in the 19th century of one-to-one media, like telegraph and telephone, and in the first half of this century of one-to-many media, like radio and TV, the last decade in par-

ticular has been characterized by the explosive development of many-to-many communication networks. Whereas the traditional communication media link sender and receiver directly, networked media have multiple cross-connections between the different channels, allowing complex sets of data from different sources to be integrated before being delivered to the receivers. For example, a newsgroup discussion on the Internet will have many active contributors as well as many people just 'listening in'. Moreover, the fact that the different 'nodes' of the digital network are controlled by computers allows sophisticated processing of the collected data, reinforcing the similarity between the network and the brain. This has led to the metaphor of the world-wide computer network as a 'global brain' [Mayer-Kress & Barczys, 1995; Russell, 1995].

In organisms, the evolution of the nervous system is characterized by a series of *metasystem transitions* producing subsequent levels of complexity or control [Turchin, 1977; Heylighen, 1995, 1991b]. The level where sensors are linked one-to-one to effectors by neural pathways or reflex arcs is called the level of *simple reflexes*. It is only on the next level of *complex reflexes*, where neural pathways are interconnected according to a fixed program, that we start recognizing a rudimentary brain. This paper will argue that the present global computer network is on the verge of undergoing similar transitions to the subsequent levels of *learning*, characterized by the automatic adaptation of connections, *thinking*, and possibly even *metarationality*. Such transitions would dramatically increase the network's power, intelligence and overall usefulness. They can be facilitated by taking the "network as brain" metaphor more seriously, turning it into a model of what a future global network might look like, and thus helping us to better design and control that future. In reference to the super-organism metaphor for society this model will be called the "super-brain".

2 The Web as an Associative Memory

The first requirement for developing a 'brain-like' global network is integration: all parts of the net should be able to communicate, using a shared protocol. At present, all existing computer networks, public and private, seem to be moving towards interconnection via the Internet and its underlying protocols [Krol, 1993]. That this network reached the critical mass where it became more attractive

* The authors are supported by the Belgian National Fund for Scientific Research (NFWO): F. Heylighen as a Senior Research Associate, J. Bollen as assistant to the project "Evolutionary construction of knowledge systems".

to link into it rather than into any competing network may be explained by the following factors: 1) its overall flexibility and robustness. 2) its public, uncontrolled character. This made it possible for anybody with a good idea to spread it quickly over the net, allowing others to use it, improve it and build on it. 3) the popularity of some of the resulting information systems, such as news-groups, Gopher, and especially the World-Wide Web (WWW), which is increasingly being used as a unified interface to all other systems. This universal acceptance is due to WWW's extremely simple, but powerful way of representing networked information: *distributed hypermedia*. It is this architecture that turns WWW into a prime candidate for the substrate of a global brain.

The distributed hypermedia paradigm is a synthesis of three ideas [Heylighen, 1994]. 1) *Hypertext* refers to the fact that WWW documents are cross-referenced by 'hotlinks': high-lighted sections or phrases in the text, which can be selected by the user, calling up an associated document with more information about the phrase's subject. Linked documents ('nodes') form a network of associations or 'web', similar to the associative memory characterizing the brain. 2) *Multimedia* means that documents can present their information in any modality or format available: formatted text, drawings, sound, photos, movies, 3-D 'virtual reality' scenes, or any combination of these. This makes it possible to choose that presentation best suited for conveying an intuitive grasp of the document's contents to the user, if desired, bypassing abstract, textual representations for more concrete, sensory equivalents. 3) *Distribution* means that linked documents can reside on different computers, maintained by different people, in different parts of the world. With good network connections, the time needed to transfer a document from another continent is not noticeably different from the time it takes to transfer a document from the neighbouring office. This makes it possible to transparently integrate information on a global scale.

Initially the Web was used for passive browsing through existing documents. The addition of 'electronic forms', however, made it possible for users to actively enter information, allowing them to create documents and query specialized computer programs anywhere on the net. At present the World-Wide Web can be likened to a huge external memory, where stored information can be retrieved either by following associative links, or by explicitly entering looked-for terms in a search engine.

2.1 The Learning Web

What it lacks, though, is the capacity to autonomously learn new information. At present, 'learning' in the Web takes place through the intermediate of people, who add documents or links to newly discovered material, using their own judgement about what is worthwhile or which documents should be linked to which other documents. However, the cognitive capacity of an individual is much too limited to get any grasp of a huge network consisting of millions of documents. Intuition is a rather poor guide for efficiently organizing the Web. The result is that the Web is mostly labyrinthine, and it is not obvious to find the information one is looking for.

A first step to make the 'Web as memory' more efficient is to let the Web itself discover the best possible organization. In the human brain knowledge and meaning develop through a process of *associative learning*: concepts that are frequently used together become more strongly connected (Hebb's rule for neural networks). It is possible to implement similar mechanisms on the Web, creating associations on the basis of the paths followed by the users through the maze of linked documents. The principle is simply that links followed by many users become 'stronger', while links that are rarely used become 'weaker'. Simple heuristics can then propose likely candidates for new links: if a user moves from A to B to C, it is probable that there exists not only an association between A and B but also between A and C (transitivity), and between B and A (symmetry). In this manner, potential new links are continuously generated, while only the ones that gather sufficient 'strength' are retained and made visible to the user. This process was tested by us in an adaptive hypertext experiment, where a web of randomly connected words self-organized into a semantic network, by learning from the link selections made by its users. [See Bollen & Heylighen, 1996, for more details about learning algorithms and experimental results].

The strength of such associative learning mechanisms is that they work *locally* (they only need to store information about documents at most two steps away), but the self-organization they produce is *global*: given enough time, documents which are an arbitrary number of steps away from each other can become directly connected if a sufficient number of users follow the connecting path. We could imagine extending this method by more sophisticated techniques, which e.g. compute a degree of similarity between documents on the basis of the words they contain, and use this to suggest similar documents as candidate links from a given document. The expected result of such associative learning processes is that documents that are likely to be used together will also be situated near to each other in the topology of 'cyberspace'.

If such learning algorithms could be generalized to the Web as a whole, the knowledge existing in the Web could become structured into a giant associative network which continuously learns from its users. Each time a new document is introduced, the links to and from it would immediately start to adapt to the pattern of its usage, and new links would appear which the author of the document never could have foreseen. Since this mechanism in a way assimilates the collective wisdom of all people consulting the Web, we can expect the result to be much more useful, extended and reliable than any indexing system generated by single individuals or groups.

3 The Thinking Web

Until now, the Web as brain metaphor has stressed its passive role as repository of knowledge, while the active search, thinking and problem-solving is done by the user, following a path through the maze of associations. A more active Web would use different information retrieval mechanisms to autonomously explore parallel documents, and deliver the combined results to the user.

A first such mechanism can be found in WAIS-style

search engines [e.g. Lycos, <http://lycos.cs.cmu.edu/>]. Here the user enters a combination of keywords that best reflect his or her query. The engine scans its index of web documents for documents containing those keywords, and scores the 'hits' for how well they match the search criteria. The best matches (e.g. containing the highest density of desired words) are proposed to the user. For example, the input of the words "pet" and "disease" might bring up documents concerning veterinary science. This method only works if the documents effectively contain the proposed keywords. However, many documents may discuss the same subject using different words (e.g. "animal" and "illness"), or use the same words to discuss different subjects (e.g. PET tomography).

Some of these problems may be overcome through a direct extension of the associative memory metaphor, the mechanism of *spreading activation* [Jones, 1986; Salton & Buckley, 1988]: activating one concept in memory activates its adjacent concepts which in turn activate their adjacent concepts. Documents about pets in an associative network are normally linked to documents about animals, and so a spread of the activation received by "pet" to "animal" may be sufficient to select all documents about the issue. This can be implemented as follows. Nodes get an initial activation value proportional to an estimate of their relevance for the query. This activation is transmitted to linked nodes. The total activation of a newly reached node is calculated as the sum of activations entering through different links, weighted by the links' strength. This process is repeated, with the activation diffusing along parallel paths, until a satisfactory solution is found (or the activation value becomes too low).

3.1 Software Agents

Until now, both search engines and spreading activation are typically implemented on single computers carrying an index of linked documents. To extend these mechanisms to the Web as a whole, we may turn to the new technology of software 'agents' [Maes, 1994]. An *agent* is a (typically small) program or script, which can travel to different places and make decisions autonomously, while representing the interests of its user.

A simple way to conceptualize the function of an agent is through the concept of *vicarious selector* [Campbell [1974]. A vicarious selector is a delegate mechanism, which explores a variety of situations and selects the most adequate ones, in anticipation of the selection that would eventually be carried out by a more direct mechanism. For example, echo-location in bats and dolphins functions through the broadcast of an acoustic signal, which is emitted blindly in all directions, but which is selectively reflected by objects (e.g. prey or obstacles). The reflections allow the bat to locate these distant objects in the dark, without need for direct contact. Similarly, an agent may be 'broadcast' over the Web, exploring different documents without a priori knowledge of where the information it is looking for will be located. The documents that fulfil the agent's selection criteria can then be 'reflected' back to the user. In that way, the user, like the bat, does not need to personally

explore all potentially important locations, while still being kept informed of where the interesting things are.

A web agent might contain a combination of possibly weighted keywords that represents its user's interest. It would evaluate the documents it encounters with respect to how well they satisfy the interest profile, and return the ones that score highest scoring to the user. Agents can moreover implement spreading activation: an agent encountering different potentially interesting directions (links) for further exploration, could replicate or divide itself into different copies, each with a fraction of the initial 'activation', depending on the strengths of the links and the score of their starting document. When different copies arrive in the same document, their activations are added in order to calculate the activation of the document. In order to avoid epidemics of virus-like agents spreading all over the network, a cut-off mechanism should be built in, so that no further copies are made below a given threshold activation, and so that the initial activation supply of an agent is limited, perhaps in proportion to the amount of computer resources the user is willing to invest in the query.

An agent's selection criteria may be explicitly introduced by the user, but they can also be learnt by the agent itself [Maes, 1994]. An agent may monitor its user's actions and try to abstract general rules from observed instances. For example, if the agent notes that many of the consulted documents contain the word "pet", it may add that word to its search criteria and suggest to the user to go and collect more documents about that topic. Learning agents and the learning Web can reinforce each others effectiveness. An agent that has gathered documents related according to its built-in or learned selection criteria can signal that to the Web, allowing the Web to create or strengthen links between these documents. Reciprocally, by creating better associations, the learning Web will facilitate the agents' search, by guiding the spread of activation or by suggesting related keywords (e.g. "animal" in addition to "pet"). Through their interaction with a shared associative web, agents can thus indirectly learn from each other, though they may also directly exchange experiences [Maes, 1994].

3.2 Solving Complex Problems

Answering queries may be further facilitated if the Web is not just associatively, but *semantically* structured, i.e. if the links can belong to distinct types with specific meanings (e.g. "is a", "has part", "has property", etc.). That would further guide searches, by restricting the number of links that need to be explored for a specific query. The resulting Web would more resemble a semantic network or knowledge-based system, capable of making 'intelligent' inferences to answer complex queries (e.g. "give me a list of all diseases that attack non-carnivorous, pet mammals"). Yet, it is important to maintain 'untyped', free associations type in order not to a priori limit the type of information that can be found in the network [Heylighen, 1991a]

We can safely assume that in the following years virtually the whole of human knowledge will be made available on the Web. If that knowledge is organized as

an associative or semantic network, 'spreading' agents should be capable to find the answer to practically any question for which an answer somewhere exists. The spreading activation mechanism allows questions that are vague, ambiguous or ill-structured: you may have a problem, but not be able to clearly formulate what it is you are looking for.

For example, imagine the following situation: your dog is regularly licking the mirror in your home. You worry whether that is just normal behavior, or perhaps a symptom of a disease. So, you try to find more information by entering the keywords "dog", "licking" and "mirror" into a web search agent. If there would be a 'mirror-licking' syndrome described in the literature about dog diseases, such a search would immediately find the relevant documents. However, that phenomenon may just be an instance of the more general phenomenon that certain animals like to touch glass surfaces. A traditional search on the above keywords would never find a description of that phenomenon, but the spread of activation in a semantically structured web would reach "animal" from "dog", "glass" from "mirror" and "touching" from "licking", selecting documents that contain all three concepts. Moreover, a smart agent would assume that documents discussing possible diseases would be more important to you than documents that just describe observed behavior, and would retrieve the former with higher priority.

This example can be generalized to the most diverse problems. Whether it has to do with how to decorate your house, how to reach a certain place, or how to combat stress: whatever the problem you have, if some knowledge about the issue exists, spreading agents should be able to find it. For the more ill-structured problems, the answer may be reached only after a number of steps. Formulating part of the problem brings up certain associations that make you or the agent reformulate the problem (e.g. excluding documents about tomography), in order to better select relevant documents. The Web will not only provide straight answers but general feedback to guide you in your efforts to get closer to the solution.

Coming back to our brain metaphor, the agents searching the Web, exploring different regions, creating new associations by the paths they follow and the selections they make, and combining the found information into a synthesis or overview, which either solves the problem or provides a starting point for a further round of reflection, seem wholly analogous to thoughts spreading and recombining over the network of associations in the brain. This would bring the Web into the metasystem level of *thinking*, which is characterized by the capability to combine concepts without the need for an a priori association between these concepts to exist in the network [Turchin, 1977; Heylighen, 1991b, 1995].

3.3 Knowledge Discovery

The next metasystem level may be called *metarationality*: the capacity to automatically create new concepts, rules and models, and thus change one's own way of thinking. This would make thinking in the Web not just quantitatively, but qualitatively different from human

thought. An intelligent Web could extend its own knowledge by the process of "knowledge discovery" or "data mining" [Fayyad & Uthurusamy, 1995]. This is based on an automatization of the mechanisms underlying scientific discovery: a set of more abstract concepts or rules is generated which summarizes the available data, and which, by induction, makes it possible to produce predictions for situations not yet observed. As a simple illustration, if after an exhaustive search it would turn out that most documented cases of dogs licking mirrors would also suffer from a specific nervous disease, a smart Web might infer that mirror-licking is a symptom of that disease and that new cases of mirror-licking dogs would be likely to suffer from that same disease, even though that rule may never have been entered in its knowledge base and been totally unknown until then.

Many different techniques are available to support such discovery of general principles, including different forms of statistical analysis, genetic algorithms, inductive learning and conceptual clustering, but these still lack integration. The controlled development of knowledge requires a unified *metamodel*: a model of how new models are created and evolve [Heylighen, 1991b]. A possible approach to develop such a metamodel might start with an analysis of the building blocks of knowledge, of the mechanisms that (re)combine building blocks to generate new knowledge systems, and of a list of selection criteria, which distinguish 'good' or 'fit' knowledge from 'unfit' knowledge [Heylighen, 1992].

4 Integrating Individuals

What is still lacking in our model in order to describe it as a super-brain, is the *integration* of individuals into a collective 'super-organism' with the thinking Web as its nervous system. This is the most controversial issue in this discussion [Heylighen & Campbell, 1995]. Yet, here too there are signs that integration, if not promoted, will at least be facilitated by the global network.

In order to most effectively use the cognitive power offered by an intelligent Web, there should be a minimal distance between the user's wishes and desires and the sending out of web-borne agents. At present, we are still using computers connected the network by phone cables, creating queries by typing in keywords in specifically selected search engines. This is quite slow and awkward when compared to the speed and flexibility with which our own brain processes thoughts. Several mechanisms can be conceived to accelerate that process.

The quick spread of wireless communication and portable devices promises the constant availability of network connections, whatever the user's location. We already mentioned multimedia interfaces, which attempt to harness the full bandwidth of 3-dimensional audio, visual and tactile perception in order to communicate information to the user's brain. The complementary technologies of speech or gesture recognition make the input of information by the user much easier. We also mentioned the learning agents, which try to anticipate the user's desires by analysing his or her actions. But even more direct communication between the human brain and the Web can be conceived.

There have already been experiments in which people managed to steer images on a computer screen simply by *thinking*: their brain waves associated with focused thoughts (such as “up”, “down”, “left” or “right”) are registered by sensors, interpreted by neural network software, and translated into commands, which are executed by the computer. Such set-ups use a two-way learning process: the neural network learns the correct interpretation of the registered brain-wave patterns, while the user, through bio-feedback, learns to focus thoughts so that they become more understandable to the computer. An even more direct approach can be found in *neural interface* research, the design of electronic chips that can be implanted in the human body and connected to nerves, so as to register neural signals [Kovacs et al., 1994]. Once these technologies have become more sophisticated, we could imagine the following scenario: at any moment a thought might form in your brain, then be translated automatically via a neural interface to an agent or thought in the external brain, continue its development by spreading activation, and come back to your own brain in a much enriched form. With a good enough interface, there should not really be a border between ‘internal’ and ‘external’ thought processes: the one would flow naturally and immediately into the other. It would suffice that you think about your dog licking mirrors to see an explanation of that behavior pop up before your mind’s eye.

4.1 Reciprocal Interaction

Interaction between internal and external brain does not need to be one-way: the Web itself might query the user. A ‘metarational’ Web would continuously check the coherency and completeness of the knowledge it contains. If it finds contradictions or gaps, it would try to situate the persons most likely to understand the issue (perhaps the authors or active users of a document), and direct their attention to the problem. An explicit formulation of the problem, possibly supported by different ‘knowledge elicitation’ techniques, is likely to be sufficient for an expert to quickly fill in the gap, using implicit knowledge which was not as yet entered into the Web [Heylighen, 1991a]. In that way, the Web would learn implicitly and explicitly from its users, while the users would learn from the Web. Simultaneously, the Web would mediate between users exchanging information or answering each other’s questions, e.g. by locating the right person at the right time, or by providing additional explanations.

In a sense, the brains of the users themselves would become nodes in the Web: stores of knowledge linked to the rest of the Web, which can be consulted by other users or by the Web itself. Eventually, the individual brains may become so strongly integrated with the Web that the Web would literally become a ‘brain of brains’: a *super-brain*. A thought might run from one user to the Web, to another user, back to the Web, and so on. Thus, billions of thoughts would develop in parallel over the super-brain, creating ever more knowledge in the process.

The question remains whether individuals would agree to be so intimately linked into a system they only partially control. On the one hand, individuals might refuse answering requests from the super-brain. On the other

hand, no one would want to miss the opportunity to use the unlimited knowledge and intelligence of the super-brain for solving one’s own problems. However, the basis of social interaction is reciprocity. People will stop answering your requests if you never answer theirs. Similarly, one could imagine that the intelligent Web would be based on the simple condition that you can use it only if you provide some knowledge in return.

In practice, such conditions may come out of the economic constraints of the ‘knowledge market’, which make that people must provide services in order to earn the resources they need to sustain their usage of other services. Presently, there is a rush of commercial organizations moving to the Web in order to attract customers. The best way to convince prospective clients to consult their documents, will be to make these documents as interesting and useful as possible. Similarly, the members of the academic community are motivated by the ‘publish or perish’ rule: they try to make their ideas as widely known as possible, and are most likely to succeed if these results are highly evaluated by their peers on the Web. Thus, we might expect a process where the users are maximally motivated both to make use of the Web’s existing resources and to add new resources to it. This will make the Web-user interaction wholly two-way, the one helping the other to become more competent.

4.2 Towards a Super-Organism?

A remaining problem is whether an integrated ‘super-brain’ will also lead to an integrated social system or ‘super-organism’. This requires not only the integration of knowledge, but the integration of the goals and values of the different users into an overarching value system steering the super-organism. In how far can the conflicting interests of all individuals and groups using the Web be reconciled and merged into a ‘global good’ for the whole of humanity? The super-brain might facilitate such an integrative process, since it is in everybody’s interest to add to the knowledge stored in the super-brain: there does not seem to be a part-whole competition [Heylighen & Campbell, 1995] between individual and super-brain. This is due to the peculiar nature of information: unlike limited, material resources, information or knowledge does not diminish in value if it is distributed or shared among different people. Thus, there is no a priori benefit in keeping a piece of information to oneself (unless this information controls access to a scarce resource!).

However, there remains the problem of intellectual property (e.g. copyright or patents): though it might be in the interest of society to immediately make all new knowledge publicly available, it is generally in the interest of the developer of that knowledge to restrict access to it, because this makes it easier to get compensation for the effort that went into developing it. An advantage of the global network is that it may automatize compensation, minimize the costs of development and transaction of knowledge, and foster competition between knowledge providers, so that the price of using a piece of knowledge developed by someone else might become so low as to make it practically free. A very large number of users paying a very small sum may still provide the de-

veloper with a sufficient reward for the effort.

As to the fair distribution of material resources over the world population, it must be noted that their value (as contrasted with intellectual resources) is steadily decreasing as a fraction of the total value of products or services. Moreover, the super-brain may facilitate the emergence of a universal ethical and political system, by promoting the development of shared ideologies that transcend national and cultural boundaries [cf. Heylighen & Campbell, 1995], and by minimizing the distance between individuals and government. However, these questions are very subtle and complex, and huge obstacles remain to any practical implementation, so that it seems impossible to make predictions at this stage.

5 Discussion

The picture we have sketched of a super-brain emerging from the global electronic network may seem more closely related to science-fiction novels than to the technical literature. Yet, the elements of this picture are methods and technologies that exist at this very moment (though sometimes still in a rudimentary form). The integrated model seems a relatively prudent extrapolation of existing developments, supported by theoretical principles from cybernetics, evolutionary theory and cognitive science. The explosive growth of the World-Wide Web, which has developed in a mere 5 years from an interesting idea into a global multimedia network, connecting dozens of millions of people, and attracting huge investments from all segments of society, shows that in the domain of information technologies the distance between concept and realization can be very short indeed.

Yet, there are the many unfulfilled promises from the 40 year history of Artificial Intelligence to remind us that problems may be much more serious than they initially appeared. It is our impression that the main obstacles hindering AI have been overcome in the present model. First, AI was dogged by the fact that intelligent behavior requires the knowledge of an enormous mass of common-sense facts and rules. The fact that millions of users in parallel add knowledge to the super-brain eliminates this bottleneck. The traditional symbolic AI paradigm moreover made the unrealistic demand that knowledge be formulated as precise, formal rules. Our view of the super-brain rather emphasizes the context-dependent, adaptive and fuzzy character of associative networks, and is thus more reminiscent of the connectionist paradigm. Finally, traditional AI tended to see knowledge as a mapping or encoding of outside reality, a philosophy that runs into a host of practical and epistemological problems [Bickhard & Terveen, 1995]. The present model, on the other hand, is constructivist or selectionist: potential new knowledge is generated autonomously by the system, while the environment of users selects what is adequate.

It will only become clear in the next few years whether these changes in approach are sufficient to overcome the technical hurdles. At this stage, we can only conclude that extensive research will be needed in order to further develop, test and implement the ideas underlying the present model for a future network.

References

- Bickhard M., Terveen L. (1995): *Foundational Issues in Artificial Intelligence & Cognitive Science* (Elsevier).
- Bollen J. & Heylighen F. (1996): "Algorithms for the Self-Organization of Distributed Multi-user Networks", in: R. Trappl (ed.) *Cybernetics and Systems '96* (this volume).
- Campbell D.T. (1974): "Evolutionary Epistemology", in: *The Philosophy of Karl Popper*, Schilpp P.A. (ed.), (Open Court Publish., La Salle, Ill.), p. 413-463.
- Fayyad U.M. & Uthurusamy R. (eds.) (1995): *Proc. 1st Int. Conference on Knowledge Discovery and Data Mining* (AAAI Press, Menlo Park, CA).
- Heylighen F. (1991a): "Design of a Hypermedia Interface Translating between Associative and Formal Representations" *Int J. Man-Machine Studies* 35, p. 491.
- Heylighen F. (1991b): "Cognitive Levels of Evolution", in: *The Cybernetics of Complex Systems*, F. Geyer (ed.), (Intersystems, Salinas, CA), p. 75-91.
- Heylighen F. (1993): "Selection Criteria for the Evolution of Knowledge", *Proc. 13th Int. Cong. on Cybernetics* (Int. Ass. of Cybernetics, Namur), p. 524-528.
- Heylighen F. (1994): "World-Wide Web: a distributed hypermedia paradigm for global networking", *Proc SHARE Europe, Spring 1994*, (Geneva), p. 355-368.
- Heylighen F. (1995): "(Meta)systems as constraints on variation", *World Futures* 45, p. 59-85.
- Heylighen F. & Campbell D.T. (1995): "Selection of Organization at the Social Level", *World Futures: the Journal of General Evolution* 45, p. 181-212.
- Jones W. P. (1986): "On the Applied Use of Human Memory Models", *International Journal of Man-Machine Studies* 25, p. 191-228.
- Kovacs G.T., Stormont C.W., Halks-Miller M. (1994): "Silicon-Substrate Microelectrode Arrays for Parallel Recording of Neural Activity in Peripheral and Cranial Nerves", *IEEE Trans. Biomed. Engin.* 41, p. 567.
- Krol E. (1993): *The Whole Internet* (O'Reilly, Sebastopol, CA).
- Maes P. (1994): "Agents that Reduce Work and Information Overload", *Comm. of the ACM* 37 (3).
- Mayer-Kress G. & C. Barczys (1995): "The Global Brain as an Emergent Structure from the Worldwide Computing Network", *The Information Society* 11 (1).
- Russell, P. (1995): *The Global Brain Awakens: Our Next Evolutionary Leap* (Miles River Press).
- Salton G. & Buckley C. (1988): "On the Use of Spreading Activation Methods in Automatic Information Retrieval", *Proc. 11th Ann. Int. ACM SIGIR Conf. on R&D in Information Retrieval* (ACM), p. 147-160.
- Stock G. (1993): *Metaman: the merging of humans and machines into a global superorganism*, (Simon & Schuster, New York).
- Turchin V. (1977): *The Phenomenon of Science. A cybernetic approach to human evolution* (Columbia University Press, New York).