

Complexity and Information Overload in Society: why increasing efficiency leads to decreasing control

Francis HEYLIGHEN

CLEA, Free University of Brussels, Pleinlaan 2, B-1050 Brussels, Belgium

<http://pcp.vub.ac.be/HEYL.html>

ABSTRACT. It is argued that social and technological evolution is characterized by ephemeralization, an accelerating increase in the efficiency of all material, energetic and informational processes. This leads to the practical disappearance of the constraints of space, time, matter and energy, and thus spectacularly increases our power to physically solve problems. However, the accompanying “lubrication” or reduction of friction in all processes creates a number of non-physical problems, characterized by the increasing instability, complexity and reach of causal networks, and therefore decreasing controllability and predictability. As a result, individuals are forced to consider more information and opportunities than they can effectively process. This information overload is made worse by “data smog”, the proliferation of low quality information because of easy publication. It leads to anxiety, stress, alienation, and potentially dangerous errors of judgment. Moreover, it holds back overall economic productivity.

1. Introduction

The explosive development of the Internet and related information and communication technologies has brought into focus the problems of information overload, and the growing speed and complexity of developments in society. People find it ever more difficult to cope with all the new information they receive, constant changes in the organizations and technologies they use, and increasingly complex and unpredictable side-effects of their actions. This leads to growing stress and anxiety, fuels various gloom and doom scenarios about the future of our planet, and may help explain the increasingly radical movements against globalization.

This paper sets out to analyse the evolutionary dynamics behind these profound societal and technological developments. The main argument will be that the technological advances that we normally would consider as progress bring with it a number of subtle, but unavoidable side effects, that make it increasingly difficult for individuals and society to control or predict further developments. Since the basic thrust of progress cannot be stopped, this means that we will have to evolve suprahuman

systems to complement our limited capacities for processing information and understanding complex systems. These systems cannot be merely technological (the famed superintelligent computers or robots), but must encompass humans as essential components. Part II of this paper will then look at how such collective systems may tackle the problem of information overload.

2. The dynamics of progress

2.1. *Ephemerization*

Progress has become a questionable concept (cf. Heylighen & Bernheim, 2000a), and many theorists have correctly pointed that what looks like improvement to one observer (e.g. creation of pest-resistant crops by genetic modification), may look like deterioration to another one (e.g. interference with nature). The interpretation obviously depends on the value system by which you judge the “goodness” of the resulting change. Yet, there is at least one form of “improvement” about which virtually everyone will agree: *doing more with less*. If you can achieve the same or more output (products, services, information, ...) while requiring less input (effort, time, resources, ...) then you have increased your power to reach your goals—whatever these goals are. This increase in power applies under all circumstances, since a reduced need for effort or resources will make you less dependent on the particular conditions under which you try to reach your goals.

Such an increase in power, productivity, or efficiency exerts a strong *selective pressure* on all evolutionary processes in society: whenever there is a competition between individuals, groups, institutions, technologies or—most generally—systems of action, then *ceteris paribus* the more productive one will win. Indeed, whatever criterion directs the competition (producing cars, providing information, selling soft drinks, making religious converts, ...), the competitor who can achieve more for the same amount of investment will have a robust advantage over the others. This means that whenever a new variant appears that is somehow more productive than its competitors, it tends to become dominant, and the others will have to follow suit, or be eliminated. Thus, as long as there is *variation* (appearance of new variants) and *selection* (elimination of the less successful variants), evolution will produce an on-going increase in productivity (Heylighen & Bernheim, 2000b). Buckminster Fuller (1969) has called this process of constantly achieving more with less *ephemerization*.

Since the development of modern science in the 17th and 18th centuries and its application to technology leading to the industrial revolution, this evolution has accelerated spectacularly. Rather than having to wait for a chance discovery, new techniques are now being developed in a *systematic* way, using the sophisticated methods for modelling and testing that characterize science. Ephemerization moreover is self-reinforcing: the greater efficiency of various systems and technologies not only leads to greater output of goods and services, but also to a faster rate of further innovation, as new ideas are generated, developed, tested and communicated with less effort. The results are staggering: culture, society and even the physical world are changing in all aspects, and this at a breakneck speed.

Some well-known examples may illustrate this accelerating change. Because of better techniques, such as irrigation, crop improvement, fertilizers, pesticides, and harvesting machines, agricultural productivity has increased spectacularly over the past two centuries: both the area of land and amount of human effort needed to produce a given amount of food has been reduced to a mere fraction of what it was. As a result, the price of food in real terms has declined with 75% over the last half century (World Resources Institute, 1998). In the same period, the fuel consumption of cars has decreased just as spectacularly, while their speed, power and comfort have increased. More generally, the average speed of transport has been increasing over the past few centuries, with the effect that people and goods need a much shorter time to reach any far-away destination. In the 16th century, Magellan's ships needed more than two years to sail around the globe. In the 19th century, Jules Verne gave a detailed account of how to travel around the world in 80 days. In 1945, a plane could do this trip in two weeks. Present-day supersonic planes need less than a day.

Without doubt, the most spectacular efficiency gains have been made in the transmission and processing of information. In pre-industrial times, people communicated over long distance by letters, carried by couriers on horseback. Assuming that an average letter contained 10,000 bytes, and that a journey took one month, we can estimate the average speed of information transmission as 0.03 bit per second. In the 19th century, with the invention of the telegraph, assuming that it takes a little over two seconds to punch in the Morse code for one character, we get a transmission rate of 3 bit per second. The first data connections between computers in the 1960's would run at speeds of 300 bit per second, another dramatic improvement. Present-day, basic modems reach some 60,000 bits per second. However, the most powerful long distance connections, using fibre optic cables, already transmit billions of bits of per second. In a mere 200 years, the speed of information transmission has increased some 100 billion times!

We see a similar explosive development of power in information processing, which follows the well-known law of Moore, according to which the speed of microprocessors doubles every 18 month, while their price halves. As a result, a single chip used in a present-day electronic toy may contain more computing power than was available in the whole world in 1960. Again, this is a beautiful illustration of ephemeralization, as more (processing) is achieved with less (time, materials).

2.2. *Reduction of friction*

The net result of the drive towards increasing efficiency is that matter, energy and information are processed and transported ever more easily throughout the social system. This can be seen as the reduction of *friction*. Normally, objects are difficult to move because friction creates a force opposing the movement. That is why car makers and airplane designers invest so much effort in aerodynamic shapes, which minimize air friction. Friction functions as a kind of resistance, which dissipates kinetic energy, and thereby slows down the movement, until complete standstill. *Noise* plays a similar role in information transmission: over imperfect lines, parts of the signal get lost on the way, until the message becomes uninterpretable. For obvious reasons, communication

engineers try to minimize noise in the same way that airplane engineers try to minimize air resistance.

Physically, friction can be seen as the force responsible for the dissipation of energy and the concomitant increase of entropy (disorder), as implied by the second law of thermodynamics. Entropy increase entails the loss of information, structure, and “free” energy, that is, energy available for performing further work. This energy must be replenished from outside sources, and therefore a system performing work requires a constant input of energy carrying resources. However, the second law only specifies that entropy must increase (or remain constant), but not how much entropy is actually produced. Different processes or systems will produce entropy to varying degrees. Ephemeralization can be seen most abstractly as a *reduction of entropy production*, meaning that inputs are processed more efficiently, with less dissipation of resources. The result is that, for a given input, a system’s output will contain more usable energy and information.

This has a fundamental consequence for cause-and-effect chains. Every process, object, or organization can be seen as an input-output system, which produces a certain output in reaction to a given input (Mesarovic & Takahara, 1975). Inputs and outputs can be material, energetic and/or informational, but they are necessarily connected by a causal relation, which maps input (cause) onto output (effect) according to a particular set of rules or dynamics that characterizes the system. Given these rules, the state of the system, and the cause or input, you can predict the effect or output. What friction affects is the *strength* of this cause-effect relationship. A high friction or high entropy relation is one in which a strong, distinct cause will produce not more than a weak, difficult to discern, effect.

Imagine a billiard-ball (system) being hit by a billiard-cue (input or cause). The kinetic energy of the hit will be transferred practically completely to the ball, making it move with a speed proportional to the momentum imparted by the cue (output or effect). Imagine now hitting with that same cue a ball made of soft clay. The kinetic energy of the impact (input) will be almost completely absorbed or dissipated by the clay, resulting in a barely perceptible movement of the ball (output). The hard, smooth billiard-ball is a low friction system, with a strong cause-effect relation. The soft, irregular ball of clay, on the other hand, is a high friction system, with a weak cause-effect relation.

Now imagine coupling different causal processes or input-output systems in a chain. The output of the first system provides the input to the next one, and so on. If all systems in the sequence would be frictionless (an extreme, unrealistic case), any input given to the first system would be transmitted without any loss of strength to all subsequent systems. If the systems have friction, though, each next output will be weaker than the previous one, until it has become so weak that it no longer has any discernible effect (see figure 1).




Fig. 1: two causal chains of systems or processes: in the top one (white systems), input is transferred to output without loss or friction; in the bottom one (grey systems), the output sequentially weakens because of friction.

Let us discuss a few examples of such causal chains. Imagine a long, straight row of billiard-balls, each ball a short distance from the next one. If you hit the first ball with your cue (cause), it will hit the second ball (effect), which will itself hit the third ball (second effect), and so on. Because of friction, energy is lost, and each next ball will move more slowly than the previous one, until the point where the ball stops before it has reached the next one in line: the causal chain has broken. If the balls, and the surface on which they move, are hard and smooth, friction will be small, and a good hit may bring a dozen balls in motion. If balls and surface are soft or irregular, on the other hand, the chain is likely to break after a single step.

For an example more relevant to society, consider food production. The initial inputs of the chain are water, nutrients and sunlight, the resources necessary to grow crops. The final output is the food consumed by people. In between there are several processing and transport stages, each accompanied by a loss of resources. For example, most of the water used for irrigation will be lost by evaporation and diffusion in the soil before it even reaches the plants. From all the plant tissue produced, a large part will be lost because it is eaten by pests, succumbs to diseases or drought, rots away during humid episodes, etc. More will be lost because of damage during harvesting and transport. Further losses occur during storage because of decay, rodents, etc. Processing the fruits or leaves to make them more tasty or edible, such as grinding, cooking, or mixing with other ingredients, will only lead to further loss. What is finally eaten by the consumer constitutes only a tiny fraction of the resources that went into the process.

As we noted above, ephemeralization has led to a spectacular reduction in these losses. In primitive agricultural systems, such as are still being used in many African countries, the output per unit of area or of water is minimal, and in bad years, hardly any produce will reach the population, leading to wide-spread famines. Modern techniques are much more efficient. For example, advanced irrigation systems bring the water via tubes directly to the root of the plant, minimizing evaporation and dissipation, and use sophisticated sensors in the leaves to monitor how much water the plant needs at any moment, so that they can supply just the amount for optimal growth. The gain

compared to traditional irrigation systems, where water runs in little canals between the fields, can be a hundredfold. Similar gains are achieved during all stages of the production and distribution process, virtually eliminating losses because of pests, decay, oxidation, etc., with the help of refrigeration, pasteurization, airtight enclosures, various conserving agents, etc.

A last example of the role of friction in causal chains will focus on information transmission. Imagine giving your neighbor a detailed account of something that happened in your neighborhood, such as an accident or a police arrest. Your neighbor tells the story to his aunt, who passes it on to her friend, who tells it to her hairdresser, and so on. It is clear that after a few of such oral, person-to-person transmissions, very few details of the original account will have been conserved, because of forgetting, omissions, simplifications, etc. Moreover, the story is likely to have accumulated a number of errors, because of misunderstandings, embellishments, exaggerations, mixing up with other stories, etc. In the end, the story is likely to be forgotten and to stop spreading, or, in the rare case that some elements have caught the public's imagination, continue to spread, but in a form that is barely recognizable compared to the original. In either case, hardly anything will remain of the initial message. A simple way to reduce such "friction" or "noise" in this chain of "Chinese whispers" is to write down the account and send it to your neighbor by electronic mail. The neighbor can then simply forward the original message to his aunt, who forwards it to her friend, and so on. Unless someone actively manipulates the text, no information will be lost, and the causal chain will extend for as long as people are willing to forward the message.

2.3. *Vanishing physical constraints*

The general effect of ephemeralization is that things that used to be scarce or difficult to obtain have become abundant. For example, in the developed countries, the problem with food is no longer scarcity but overabundance, as people need to limit their consumption of calories in order to avoid overweight. Even in the poorest countries, the percentage of people that are undernourished is constantly decreasing (Goklany, 2000; Simon, 1995). More generally, the trend is clearly visible in the spectacular growth in wealth, usually measured as GDP per capita, since the beginning of the 19th century (Goklany, 2000). The ever increasing productivity not only results in people earning more, but in them working less hours to achieve this wealth. Moreover, this economic development is typically accompanied by a general increase in the factors that underly overall quality of life: health, safety, education, democracy and freedom (Heylighen & Bernheim, 2000a; Simon, 1995; Goklany, 2000).

This is of course not to say that we live in the best of possible worlds. Many things are still much less abundant than we would like them to be, and although increasing productivity leads to an ever more efficient use of natural resources (Heylighen & Bernheim, 2000a), ecologists have rightly pointed out that our present usage of many resources is unsustainable. The focus of this paper, though, is not on the remaining scarcities and wastages, which ephemeralization hopefully will sooner or later eradicate, but on a wholly new category of problems created by the emergence of

“hyperefficient” processes. To get there, we first need to understand more fundamentally how ephemeralization affects the dynamics of society.

In practice, most of the physical constraints that used to govern space, time, matter, energy and information have vanished. In the Western world we can basically get as many material goods and as much information as we need, and this for a negligible investment in time and energy. (Of course, you can always *desire* more than you may need or be able to get). Moreover, distance as a factor has become largely irrelevant, as it costs hardly more effort to get goods, services of information from thousands of miles away, than from a neighboring quarter. This is the true meaning of *globalization*: the observation that social, economical and cultural processes no longer are impeded by geographical borders or distances, but cover the world as a whole. This is most clear on the Internet, where you can exchange information virtually instantaneously, without being aware whether your correspondent is situated around the corner, or on the other side of the planet. This practical disappearance of distance constraints has been referred to as *the death of distance* (Cairncross, 2001), or *the end of geography* (O'Brien, 1992).

Similarly, most of the constraints of duration have disappeared: apart from large-scale developments (such as building a house), most of the things an individual might need can be gotten in an interval of seconds (information, communication) to hours (most consumer goods and services). (in the Middle Ages, on the other hand, most of these commodities might have demanded months to acquire them, if available at all). Just imagine that you sit at your desk and suddenly start feeling hungry: a single phone call or web visit is sufficient to order a pizza, which will be delivered at your door 15 minutes later. The result may be called the *real-time society* (Larsson & Lundberg, 1998): soon, all your wishes will be fulfilled virtually instantaneously, with a negligible waiting time.

Energy too is no longer a real constraint on the individual level: practically any system that we might need to produce some work or heat can just be plugged into the ubiquitous electricity network, to get all the energy it needs, for a price that is a mere fraction of our income. Finally, matter too becomes less and less of a constraint in any practical problem-solving. The raw material out of which a good is made (e.g. steel, plastic, aluminum) contributes ever less to the value of that good. In spite of dire warnings about the exhaustion of limited reserves, the real price of physical resources (e.g. copper, tin, coal, ...) has been constantly decreasing over the past century (Simon, 1995), and has become insignificant as a fraction of the income we spend on consumption. This has led to a post-industrial economy that is mostly based on the exchange of immaterial resources, such as knowledge, organization and human attention. It is on this level, as we will see, that fundamentally new, “cybernetic” constraints are appearing.

3. Losing control

The elimination of friction has great benefits for facilitating desired processes. However, it can be dangerous when there is a risk for unwanted processes. For example, ice as a surface produces much less friction than concrete. That is why you can reach higher speeds and sustain them for a longer time when skating than when running. However,

walking on ice is much more difficult and potentially dangerous than walking on concrete: once you start slipping there is very little to stop the movement getting out of control.

In a similar way, ephemeralization smoothens or *lubricates* the machinery of society. Movements of matter and information run more freely, with very little loss or resistance. But this applies to unwanted movements too. It has become much easier to distribute weapons, bombs, drugs or poisonous materials, or for criminals or terrorists to coordinate their activities across borders. Let us go back to our two examples of food production and story spreading. Imagine that at some stage of the food production process, there is a contamination with a chemical or biological toxin. Because of the great efficiency of the chain, the poison can spread far and wide, in a very short time, and with practically no dilution or dissipation. Some real-world examples are the case of contaminated olive oil in Spain, which led to several deaths, and the dioxin scandal in Belgium (Bernard et al., 1998), which created a national scare and led many countries to stop importing Belgian food. In the latter case, the contamination was traced back to vegetable oil being recycled into animal feed, but that probably got mixed up with some motor oil containing the cancer-producing chemicals dioxin and PCBs. The feed was given to chickens, whose meat and eggs were sold and consumed all around the country before the contamination was discovered.

Similar dangers characterize friction-free information transmission. A false, misleading announcement will be forwarded just as efficiently through email as a true account. Well-known examples are the many warnings about non-existing computer viruses (Symantec, 2002), that are forwarded from person to person, and that create unnecessary panic in millions of ill-informed computer users.

3.1. *Runaway processes*

Such problems are compounded when the causal signal is not just maintained but amplified. Since there is always some loss through friction, if you want to keep a signal intact, you need to regularly amplify it, by adding external energy to the process. This principle underlies information transmission over various electrical and optical cables.

In this case, the intention is to keep the intensity of the signal constant, but the injection of external resources into the causal chain can also make the intensity grow at each stage of the sequence. This is typical of chain reactions or snowballing processes, where every increase in intensity produces further increases. This leads to an explosive growth, which only stops when all available resources have been exhausted. In such a process with *positive feedback*, the input of resources is the bottleneck, determining how fast and how large the snowball will grow. Ephemeralization, by making resources more available, widens this bottleneck. This makes explosive developments much more common, and increases the probability that the system would get out of control.

Let us again consider some examples. In a world where people and goods can travel easily, infectious agents such as viruses and bacteria too can travel more easily. Since one person can infect several others, there is a tendency for amplification: at each causal step, the number of infections increases. The resources the infection needs are people, and physical exchanges between people. In a densely populated world, where there are

lots of direct and indirect contacts between people, the spread of the infection is potentially explosive.

Computer viruses are a more modern variant of the same principle: the easier and faster the exchange of information between computers, the more explosive their spread. Suppose that every day a virus infects another computer from an already infected computer. The doubling period of the process is one day, since after that interval there will be two infected computers instead of one. After two days, there will be four, after three days eight, and so on. If nothing is done to stop the spreading, after ten days there will be about a thousand infected computers, and after twenty days a million. Thousand infected computers seems about the stage where a directed intervention becomes likely, since there are enough cases to get a clear diagnosis of the problem, but the problem has not yet gotten out of hand. At the million infections stage it is likely that large parts of the network have shut down and become incapable to react. Ten days seems like a short, but reasonable time to set up an intervention. But imagine now that the computer network would be more efficient, transmitting information, including viruses, at a much higher rate. If it would be twice as fast, the million infections mark would be reached after ten days instead of twenty, making a successful intervention unlikely. If the net would be ten times as fast, two days would be sufficient to infect a million computers, and any concerted action would appear impossible.

Another example is the 1987 "Black Wednesday" collapse of stock prices, which was due not so much to the state of the economy, but to the new phenomenon of computer trading. Specialised computer programs would monitor the prices of different stocks. If prices fell below a certain value, the computer was programmed to offer the shares for sale, before they would lose even more value. But the more shares were on sale, the lower their prices became, triggering yet more selling. This development reinforced itself, producing ever more selling for ever lower prices. Since networked computers would get data about prices from all around the world, and react immediately, this led to an extremely fast, world-wide collapse of the share indexes.

The positive feedback process, where buying triggers more buying (causing a "boom") and selling triggers more selling (causing a "bust"), was not new. What was new, was the absence of friction. Normally, a speculator would need time to get information about the price, and to give the order to buy or sell. This very much slowed down the process. It would provide a delay during which speculators would have time to assess whether the change in price was caused by a change in the intrinsic value of the stocks, or just by a temporary fluctuation. In the latter case, they could decide to buy undervalued stocks, under the assumption that the downward movement would soon be reversed. In the computer controlled market, there simply was no time for such corrective action. Shares would lose most of their value before anybody could assess what they were really worth.

The investigation of the Black Wednesday crash concluded with the recommendation that delays should be built into the computer trading programs. Another such proposal to artificially add friction is the tax proposed by the economist James Tobin to discourage currency speculation and other overly fast international money transfers that increase the volatility of the global economy. Such attempts seems

paradoxical in a world where increased speed is the highest good. Economists are already planning for the coming "superliquid", frictionless economy, where the slightest increase in demand would be immediately satisfied by a corresponding increase in supply. Yet the above examples remind us that we need additional controls when implementing frictionless mechanisms (cf. Heylighen, 2002).

3.2. *Increasing complexity of causal networks*

Ephemeralization not only lengthens causal sequences, it increases the number of effects that a cause produces in parallel. An event has generally more than one effect simultaneously. A poor harvest may lead to local famine, but it may at the same time lead to increased food prices, higher imports, political unrest, and an increase in the stock values of food producers. Each of these effects will have several further effects, which in turn produce many more effects, some of them feeding back into the original causes, thus resulting in an increasingly intricate network of interactions. Obviously, reduced friction will increase the overall number of factors affected by the initial cause, and therefore the complexity of this network (see figure 2).




Fig. 2: Increasing complexity of causal networks: in a hypothetical high friction case (left), effects die down after two steps, affecting only a small number of factors; in the low friction case (right), the number of effects grows roughly exponentially with the number of steps that the causal chain extends. Although this is not depicted, effects may moreover feed back into their causes, further complicating the picture.

The phenomenon may be illustrated by going back to our billiard-ball example. Imagine that the billiard-balls aren't neatly aligned in a row but scattered over the billiard-table. The ball you hit with your cue will now hit several balls rather than one. Each of those hits several further balls, and so on. Some of those will rebound against the sides of the table and come back to hit the ball that initially started the movement. In the end, balls coming from all possible directions are colliding, resulting in an overall chaotic movement. The lower the friction, the more collisions there will be and the

longer it will take for the movement to settle down. The resulting configuration will be unrecognizable, no ball remaining in the place it was.

The effect on society of this extension of causal networks is a greater *interdependence* of various subsystems and processes. Any action will have an increasing number of unanticipated or unintended consequences: *side-effects*. This entails a greater difficulty to predict, and therefore control, the overall effects of any particular event or process. The reduction of friction in causal *chains* merely increases the speed of the process, the number of subsequent effects in the chain, and the risk of snowballing. This reduces controllability but not necessarily predictability: as long as the cause-effect relationships are known, it is relatively easy to determine when and in how far a particular process will affect another one. The reduction of friction in causal *networks*, however, makes prediction ever more difficult, since the number of factors that need to be taken into account to determine any one outcome explodes, while the myriad interactions between those factors are likely to make the overall process ever more *chaotic*, i.e. sensitive to the smallest changes in initial conditions.

It must further be noted that the evolutionary dynamic underlying ephemeralization not only increases the complexity of interactions, but also the complexity of the overall system because it promotes the differentiation and integration of subsystems (Heylighen, 1999, 2002): ever more diverse and specialized organizations emerge, that become ever more dependent on other organizations for their inputs (suppliers) and outputs (clients). The net result is that any phenomenon, system or process in society becomes more difficult to analyze, model, predict and control.

Yet, individuals and organizations must be able to predict and control the events in which they take part at least to some degree if they wish to survive and reach their goals. To compensate for the loss of predictability, this means that they will have to gather more extensive information about all the different factors and interactions that may directly or indirectly affect their situation. This additional information may provide early warnings for unanticipated problems or deviations, and may function as building blocks for the construction of more sophisticated models. For example, to achieve any degree of reliability in predicting the utterly complex and chaotic weather system, meteorologists need very detailed and extensive data about temperature, barometric pressure, cloud cover, wind speed, etc., in all parts of the world.

Happily, ephemeralization has made the collection and processing of information much easier. For example, meteorologists can simulate weather dynamics on supercomputers, using fine-grained data provided in real-time through satellite observation. However, this spectacular improvement in capacity has largely ignored the fundamental bottleneck: the human decision-maker.

4. Information overload

Information and communication technology has made information abundant: thanks to the Internet you can basically get any information you might desire in seconds. During most of history, information was a scarce resource that was of the greatest value to the small elite that had access to it (Shenk, 1997). Enormous effort would be spent in

copying and transferring the little data available, with armies of monks toiling years in the copying by hand of the few available books, and armies of couriers relaying messages from one part of the kingdom to another. Nowadays, it rather seems that we get much more information than we desire, as we are inundated by an ever growing amount of email messages, internal reports, faxes, phone calls, newspapers, magazine articles, webpages, TV broadcasts, and radio programs.

Part of the problem is that ephemeralization has made the retrieval, production and distribution of information infinitely easier than in earlier periods, practically eliminating the cost of publication. This has reduced the natural selection processes which would otherwise have kept all but the most important information from being transmitted. For example, given the ease and low cost of Internet use, anyone can type in a message in 5 minutes, and with one command make it available on a website or send it by email to thousands of people. The result is an explosion in irrelevant, unclear and simply erroneous data fragments. This overabundance of low quality information has been called *data smog* by Shenk (1997). The worst variety are the junk email messages, commonly called "spam", carrying commercial publicity, chain letters, crackpot opinions, or scams, that are distributed automatically to millions of email addresses which have been harvested from a variety of places. Experienced Internet users will undoubtedly have noticed how the amount of spam they receive increases almost month by month, and on-going ephemeralization can only point to a continuation of that trend. It is clear that sooner or later radical measures will have to be taken to limit this shower of unwanted information.

But not only irrelevant or unwanted messages increase in number. The messages we get from friends, family members, bosses, collaborators or members of the same interest groups also shows a relentless increase. Though most of those may not be strictly necessary, they are potentially relevant to the way we live. The same applies to the ever growing amount of information that reaches us via the mass media, via television news, articles or reports. Perhaps a rebellion in some far-away African country may trigger an influx of refugees to your country, or boost the value of your shares in copper mines. The discovery that the use of a certain type of vitamin decreases the chances of getting colon cancer may save your life. The potential relevance of apparently far-away events to your own situation is moreover increasing because of the friction reduction we discussed before. For example, after the terrorist attack on the New York World Trade Center, billions of people world-wide suddenly felt the need to be informed about the arcane politics of the small nation of Afghanistan. If causal chains travel more far and wide, we also need to keep more broadly informed about the remote causes of those chains. So, although we may wish to limit our intake of information, ephemeralization forces us to pay attention to ever more data.

The problem is that people have clear limits in the amount of information they can process. To use Simon's (1972; Simon et al., 1992) well-known phrase, they have *bounded rationality*. The best known limitation is the "magical number" that governs short-term memory: the psychologist Miller (1957) has shown that people can only keep some seven items at once in their working memory. Moreover, there are clear limits on the speed with which the brain can process items in working memory. One estimate

(Csikszentmihalyi, 1990) therefore calculates its maximum processing capacity as 126 bits per second. While these results are paltry in comparison with present-day computers, it is clear that the brain still has a number of deep powers that computers lack, and therefore it is very difficult to calculate true capacity. Long-term memory is much more powerful and can store millions of concepts, although it is short-term memory that we use to think, decide, and solve problems in real-time. Still, even the capacities of long-term memory pale in comparison to the billions of documents that are already available on the web.

The practical effect of these limitations is that at a certain stage, people will be confronted with more information than they effectively can process: this situation we may call *information overload* (Berghel, 1997; Kirsh, 2000). This means that part of that information will be ignored, forgotten, distorted or otherwise lost. The problem is that we cannot say which part that is: the only way to tell is to compare the information that eventually got processed with the information that we initially got subjected to; yet by our assumption, the latter is too large for us to consider. By definition, the separation between the information that is processed and the one that is not, happens in an only partially conscious, haphazard way, as consciously rejecting information requires processing it. The result is that an individual in a situation of information overload will not only miss out on potentially important information, but moreover be aware that something is missing, while not knowing precisely what is missing and thus feeling a loss of control (Wurman, 1990). Frantic efforts to compensate for the missing data by garnering additional information are self-defeating, as they will merely further bring into focus the intrinsic limitations of the human capacity for cognition.

Although from a theoretical point of view the existence of such hard-wired limitations is obvious, in practical situations it is very difficult to estimate precisely how much information a given individual can assimilate. Incorrect estimates will lead either to misplaced confidence, as when a person believes that he or she is well-informed but actually has overlooked some crucial pieces of the puzzle, or to feelings of guilt or shame, as when people know that they cannot cope with the situation, but believe it is because they haven't researched it hard enough or are too stupid. As Csikszentmihalyi (1990) has studied extensively, a basic condition for well-being is that the challenges of the situation match a person's skills; whenever the challenges become higher than the skills, well-being is replaced by anxiety and loss of control.

The net result on the individual is increasing *stress*, and its concomitant physical, psychological, and social problems (cf. Heylighen & Bernheim, 2000b; Wurman, 1990). The longer people are subjected to information overload, the more negative its effects on physical and mental well-being. A world-wide survey (Waddington, 1996) found that two thirds of managers suffer from increased tension and one third from ill health because of information overload. The psychologist David Lewis, who analysed these findings, proposed the term "Information Fatigue Syndrome" to describe the resulting symptoms. They include anxiety, poor decision-making, difficulties in memorizing and remembering, reduced attention span, reduced work satisfaction and strained relations with collaborators (Waddington, 1996; Shenk, 1997; Wurman, 1990). In certain individuals, such enduring loss of control may further lead to helplessness, depression,

and the increasingly common “burn-out” syndrome. In the population at large, the resulting experience of ever increasing complexity may lead to *alienation*, a feeling of powerlessness, meaninglessness and lack of understanding (Geyer, 1992), and to the loss of confidence in institutions, such as governments, police, law, companies and churches, that are seen to fail in their function of controlling these complexities (Heylighen & Bernheim, 2000b; Nye et al., 1997).

4.1. *Opportunity overload*

If we consider individuals as goal-seeking, cybernetic systems (cf. Heylighen, 2002), then processing incoming information (perception) is only one half of the story. The purpose of the processing is to interpret the information with respect to the individual’s goals and values, so that this interpretation of the perceived situation can help us to reach a decision about which action to take. This preparation of an (outgoing) action is the other half. Ephemeralization has boosted not only the availability of information but our capacity for action. We have ever more numerous and more powerful tools, support systems, services and products at our disposal. A simple click on a button may be sufficient to set in motion a complex chain of events, starting in a different continent, and leading to the delivery of a rare, desired item at your doorstep. Practically any good—from cars to flowers, antiques, books and pizzas—, or service—from medical treatment to adventure travel, carwash, religious support, or house cleaning—can be ordered through a single phone call or electronic form submission.

The problem is not so much getting the desired action, as deciding which out of millions of available possibilities to choose. Would you prefer a cruise through the Caribbean to a visit of the pyramids, a romantic weekend in Paris, a trekking in the Rocky Mountains, or a guided tour along the Great Wall of China? To tackle your recurrent feelings of fatigue, would you try this new herbal tea, this antioxidant cocktail, that water-based therapy, this relaxation technique, or that electromagnetic stimulation device? Whatever the type of action you are considering, the number of possibilities has in practice become endless. This makes it ever more difficult to make a motivated choice among the alternatives. We may call this the problem of *opportunity overload*.

Opportunity overload shares many properties with information overload. First, opportunity overload is to some degree frivolous, as lots of possibilities or opportunities proposed to an individual actually offer little value, similar to data smog or spam. They merely cloud the picture, and might as well be left out altogether. Second, however, there is no simple way to separate the wheat from the chaff, and certain differences between options that appear superficial may turn out to be vital. For example, the electromagnetic stimulation device may actually turn out to be dangerous, while the antioxidant cocktail may not only give you more energy but improve your overall health and life-expectancy. Because of ephemeralization, the potential power of actions, whether to the good or to the bad, has tremendously increased. Therefore, making the right decision has become more important. But bounded rationality means that our capacity for decision-making is limited, and we cannot systematically consider *all* options. Again, the result is stress, and a constant fear that you are making the wrong

decision, or that you have failed to explore the one option that would optimally resolve all your problems.

4.2. *The productivity paradox*

While it seems pretty obvious that information overload produces stress and loss of control, it may have another, more paradoxical side effect. Ephemeralization was defined as an on-going and accelerating increase in the productivity of all processes. Yet, if we look at the economic statistics of the last three decades, overall productivity of labor and capital seems to have been increasing only slowly and haphazardly. While information technology started invading the office around 1980, any positive effects on productivity only became noticeable from 1995 on (Oliner and Sichel, 2000; Triplett, 1999). Until then, the situation could be summarized by Solow's (1987) famous quote: "You can see the computer age everywhere but in the productivity statistics". A plausible explanation for such a lag, as proposed by Perez (1983), is social and institutional *inertia*: individuals and organizations need years, if not decades, to adapt to a new technology and to learn to use it productively. This interpretation is supported by David's (1990) case study of electrical power, which needed some 20 years before its introduction (around 1900) led to significant productivity increases.

Information overload would appear to be a fundamental contributor to this inertia. New technologies increase the number of options to be considered and the amount of information that needs to be assimilated. The broader and more powerful the technology, the larger the increase. The required effort to study these new data with its accompanying stress (e.g. in the form of the information fatigue syndrome mentioned above) will obviously reduce the productivity of employees, and increase the probability of errors and ineffectual use.

However, there is an even more direct way in which complexity and information overload hold back productivity. Consider the introduction of a key technology, such as electricity, the personal computer, or the Internet. The number of possible applications or uses of such technology is extremely large, not to say infinite. Some of these applications will lead to a significant productivity increase, while others may produce only a marginal or zero increase, or even decrease. However, for the new users these potential effects are everything but obvious. The complexity, interdependence and non-linearity of processes will lead to many unexpected, counterintuitive and sometimes even vicious side-effects.

For example, putting your email address on your web home page would appear like an obvious method to make it easier for friends, colleagues, and customers to contact you. Yet, such public email addresses are now being harvested by the millions by web robots, in order to be sold to would-be spammers, who will then inundate you with junk mail, thus drowning out the messages that you really care to receive. A practical solution, discovered independently by many experienced Internet users, is to "camouflage" your precise email address, so that it can be recognized by a person wanting to contact you, but not by an automatic harvesting program.

Another argument to explain why immediate productivity gains are elusive is the *principle of suboptimization* (Machol, 1965): optimizing a subsystem or subprocess

does not in general optimize the overall system. The reason is that a complex system (e.g. a company) is more than the sum of its parts. Because of non-linear interactions, increasing the output of one part does not in general lead to a proportional increase in the overall output. A classic example is the bottleneck phenomenon: if a process depends on one scarce, but critical resource, then efficiency gains in the use of other resources won't have any effect on the total amount of production.

To find the most productive use of a new technology, avoiding all unexpectedly negative side-effects, you should systematically study *all* possible options and *all* their direct and indirect consequences. Because of bounded rationality, however, this is obviously impossible. Instead, improvements are typically made by trial-and-error, where first the most obvious applications are tried out, and users slowly become aware of their positive and negative effects. On the basis of that experience, they learn how to avoid clearly unproductive applications. They then try out minor variations or improvements on the apparent “best practices”, and perhaps carry out a few more daring experiments that may open wholly new domains of application. They thus gradually discover more productive ways to use this same technology, ending up with initially unimagined uses.

The more revolutionary the technology, the longer this learning or adaptation process will take, but the farther it will go. For example, a refrigerator was an innovation with pretty obvious applications, doing little more than replacing the existing ice cellars. Requiring little learning or information processing from its users, it increased productivity almost immediately—albeit in a restricted domain. A computer, on the other hand, initially seemed to be little more than a faster calculator or file cabinet. With the new applications of word processing and desktop publishing, though, for many it became the equivalent of their personal printing press. With multimedia, it became moreover a personal art and movie studio. With the web, it became a medium for electronic publishing, obviating the need to print copies of a document. For companies, the web itself initially seemed to be a tool merely for publishing electronic publicity leaflets and catalogs. At the moment, its great potential for electronic commerce and knowledge exchanges between companies, customers and employees is being explored. Much more innovative applications are still waiting in the wings, some of which will be discussed in the subsequent paper (Heylighen, submitted).

All of these novel applications potentially increase economic productivity, in a broad variety of domains. Yet, it is clear that we are only scratching the surface and that the productivity contributions of present-day information technologies are just a fraction of what they could be. The faster growth associated with the “new economy” is probably a mere first indication of what is possible. Yet, the “irrational exuberance” of the market reaction (the dotcom boom) and the resulting recession reminds us that non-linear side-effects and socio-institutional inertia are not that easily overcome.

The productivity paradox in most general terms then is that revolutionary technologies developed to increase productivity initially fail in this purpose, because their introduction adds to complexity and information overload, obscuring their potentially most productive applications while increasing the workload and confusion of

the individuals that have to learn how to use them. This very much slows down economic growth, and progress in general.

The fact that our present, technologically very advanced society still has not succeeded to eliminate such ancient scourges like hunger, poverty, illiteracy and epidemics may well be largely due to the resulting lag in productivity. For example, given the efficiency of present agricultural technologies there is no reason why any part of the world should ever lack food. In regions like the EU and the USA the problem is rather overcapacity, forcing ever more farmers to give up their jobs and farming land to be abandoned. Yet, these applications take decades to diffuse to the less developed countries, because of the steep learning curve they impose on a poorly educated population and ill-prepared institutions, and because of various bottlenecks and other non-linear dependencies (e.g. it doesn't help to increase the productivity of agriculture if the crops are left rotting in the field because of lacking transport or storage facilities).

5. Conclusion

Ephemerization, the ongoing increase in efficiency or productivity of all processes involving matter, energy and information, is the most basic manifestation of technological and organizational advance. It is the motor behind economic and social progress. Its general effect is to reduce friction, i.e. the loss of resources, time and information, in all socio-technical systems. Although this vastly increases our power to tackle existing problems, eliminating most physical constraints, it creates a number of novel, non-physical problems.

Most fundamentally, these are caused by the increasing reach and complexity of causal networks, making prediction and control ever more difficult. More concretely, they boil down to the increasing difficulty of making decisions, that is, selecting a subset of options from a much larger set of opportunities for action or potentially relevant pieces of information. While ephemerization constantly increases the number of items that need to be paid attention to, there is an upper bound to the human capacity for the needed processing. The result is an ever increasing gap between what an individual is cognitively able to do, and what (s)he perceives necessary to do. This leads to both subjective frustration, where people feel anxious or guilty because they think they may have missed essential elements, and objective failure, where wrong decisions are made because not enough information was taken into account.

On the level of society this produces stress and alienation, instability, problems snowballing out of control, and an overall economic growth much lower than what could be expected from the increased efficiency of individual processes. Since ephemerization cannot be stopped, it is clear that these problems will only worsen, unless fundamental solutions are introduced. The subsequent paper in this series (Heylighen, submitted) will discuss the most basic approaches for coping, and argue that only a suprahuman system, integrating both people and information technologies, can effectively solve the problem.

References

- Berghel H. (1997): "Cyberspace 2000: dealing with information overload", *Communications of the ACM* 40 (2), p. 19-24.
- Bernard, A. C. Hermans, F. Broeckaert, G. De Poorter, A. De Cock, & G. Houins (1999): "Food contamination by PCBs and dioxins: an isolated episode in Belgium is unlikely to have affected public health", *Nature*, 401, 16 Sept., p. 231-232.
- Buckminster Fuller, R. (1969) *Utopia or Oblivion*. Bantam, New York.
- Cairncross F. (2001): *The Death of Distance: How the Communications Revolution Will Change Our Lives*, Harvard Business School Press,
- Csikszentmihalyi M. (1990) *Flow: The Psychology of Optimal Experience*. New York: Harper and Row
- David P. A. (1990) "The Dynamo and the Computer: an Historical Perspective on the Modern Productivity Paradox." *American Economic Review*. 80, p. 355
- Geyer R. F. (1992): "Alienation in community and society: effects of increasing environmental complexity", *Kybernetes* 21(2), p. 33-49.
- Goklany, I. M. (2000), "Economic Growth and the State of Humanity", PERC Policy Series, <http://www.perc.org/ps21.pdf>.
- Heylighen F. (1999): "The Growth of Structural and Functional Complexity during Evolution", in: F. Heylighen, J. Bollen & A. Riegler (eds.) (1999): *The Evolution of Complexity* (Kluwer, Dordrecht).
- Heylighen F. & Bernheim J.(2000a): "Global Progress I: empirical evidence for increasing quality of life", *Journal of Happiness Studies* 1 (3), p. 323-349,
- Heylighen F. & Bernheim J. (2000b): "Global Progress II: evolutionary mechanisms and their side-effects", *Journal of Happiness Studies* 1(3), p. 351-374
- Heylighen F. (2002): "The Global Superorganism: an evolutionary-cybernetic model of the emerging network society", *Journal of Social and Evolutionary Systems*
- Heylighen F. (submitted): "Tackling Complexity and Information Overload: intelligence amplification, attention economy and the global brain"
- Kahneman (1973) *Attention and Effort* (Prentice-Hall, Englewood Cliffs NJ)
- Kirsh D. (2000): A few thoughts on cognitive overload, *Intellectica*.
- Larsson, M. & Lundberg, D. (1998): *The Transparent Market: Management Challenges in the Electronic Age*. New York: St. Martin's Press, .
- Machol, Robert E. (1965): *System Engineering Handbook*. McGraw-Hill, .
- Mesarovic, MD, & Takahara, Y: (1975) *General Systems Theory: Mathematical Foundations* Academic Press, New York
- Miller, G.A. (1957), "The magical number seven plus or minus two: some limits in our capacity for processing information", *Psychological Review*, p. 63-97
- Nye J.S, Zelikow D., et al. (eds.) (1997): *Why People Don't Trust Government*, (Harvard University Press, Cambridge MA).
- O'Brien R. (1992): *Global Financial Integration: The End of Geography*, (London: Pinter Publishers).
- Oliner, S.D. & Sichel, D.E. (2000) The Resurgence of Growth in the late 1990s: Is information Technology the Story?, *Journal of Economic Perspectives*, 14:4, p. 3-22
- Perez, C. (1983): "Structural Change and Assimilation of the New Technologies in the Economic and Social Systems", *Futures*, 15(5), p. 357-375.
- Waddington P. (1996): *Dying for Information? An Investigation into the Effects of Information Overload in the UK and Worldwide* (Reuters Business Information report, London).

- Shenk D. (1997): *Data Smog: Surviving the Information Glut* (Harper, San Francisco).
- Simon H. (1972): "Theories of Bounded Rationality," in *Decision and Organization*, R. Radner (ed.), North Holland.
- Simon H., Egidi, M., Marris, R., and Viale, R. (1992). *Economics, Bounded Rationality and the Cognitive Revolution*. Edward Elgar.
- Simon J. L. (ed.) (1995) *The State of Humanity*. Blackwell, Oxford.
- Solow, R. M. 1987. We'd better watch out. *New York Times Book Review* (July 12), p. 36.
- Symantec (2002): Hoaxes <<http://www.symantec.com/avcenter/hoax.html>>
- Triplett J. E. (1999): The Solow productivity paradox: what do computers do to productivity? *Canadian Journal of Economics* 32: 2, p. 309-334
- Wurman R.S. (1990): *Information Anxiety* (Bantam).

Tackling Complexity and Information Overload : intelligence amplification, attention economy and the global brain

Francis HEYLIGHEN

CLEA, Free University of Brussels, Pleinlaan 2, B-1050 Brussels, Belgium

<http://pespmc1.vub.ac.be/HEYL.html>

ABSTRACT. Because of accelerating technological evolution, society is confronted with more information and complexity than it can effectively process. Since this evolution cannot be stopped, this means that we will have to drastically increase our capacity for processing. This paper considers three subsystems whose capacity can be increased: individual intelligence, computer intelligence, and the transaction mechanisms that allocate processing capacity. It is concluded that each of them can be substantially improved, by measures such as education, training, drugs, information hygiene, and an economy of attention. Yet, the overall problem can be tackled only by integrating these partial approaches into an encompassing, collective or distributed, cognitive system: the “global brain”. The functioning of this system is illustrated by some concrete applications: an intelligent transport system, the routing of email messages, and the learning of new links by the web. The system as a whole must self-organize, using principles inspired by ant algorithms, Hebbian learning, and the law of supply and demand. By searching for optimizations at both the individual and the collective level, it will provide a ubiquitous, flexible decision-support that can in principle solve all problems caused by complexity and information overload.

1. Introduction

It does not need to be argued that the people in present-day society suffer from increasing complexity and information overload. The preceding paper (Heylighen, 2002) discussed the causes and effects of this phenomenon.

The fundamental cause was argued to be the process of *ephemeralization*: the ever increasing productivity or efficiency of all processes brought about by technological innovation. Ephemeralization facilitates the transfer and processing of all matter, energy and information and thus lubricates all systems in society. This increases our power to solve problems and to reach our goals. However, it has the fundamental side-effect of increasing the speed and strength of cause-effect relationships, so that any given event produces more numerous and potentially more devastating effects in a given lapse of time. This makes it much harder to predict and control the effects of our actions, while at the same time increasing the number of actions that can potentially be executed.

This on-going increase in the complexity of interaction sequences present itself to the individual as the problem of *information overload*: the amount of data and of options that an individual must take into account in order to make an adequate decision is larger than that individual's capacity for processing the information. This is worsened by *data smog* (Shenk, 1997): the proliferation of low quality or "parasitic" information made possible by technologies for the easy and inexpensive distribution of information. The result is anxiety and stress on the one hand, and potentially disastrous errors of judgment on the other hand. Society too is confronted with overload, as its institutions are unable to effectively cope with the massive amount of change to which they are subjected and decisions that need to be made. This leads to an increased risk of problems getting out of control, especially if these involve a positive feedback or vicious cycle. It also brings about a "productivity paradox", i.e. a generalized lag in the productive application of new technologies, as their introduction creates interactions and side-effects that are too complex to be rationally anticipated.

These problems are severe, imposing a great stress on society. Continuing ephemeralization can apparently only make them worse, increasing the probability of catastrophic outcomes. Yet, the present paper wishes to argue for a more optimistic view, by examining some mechanisms that would allow us to cope. We will first consider separate approaches on the level of the individual, society, and supporting technology. We will then argue that the only true solution must involve the synergetic use of these three components, in the form of an emergent collective intelligence, or "global brain". This will be illustrated by an-depth examination of some representative applications of such a distributed information system.

2. Increasing human intelligence

An obvious way to reduce the gap between the amount of information that must be processed and the one that can be processed is to increase our capacity for processing. There exist plenty of tried and true methods to increase knowledge, wisdom and intelligence, that is, the cognitive basis for making good decisions. The benefits of *education* do not need to be argued. The higher the level of education people reach, typically the more complex the problems they can tackle, and the better their understanding of the intricate causal networks in which they participate. This makes it more likely that an educated person will consider and understand the relevant information and make a wise decision.

Happily, ephemeralization is accompanied by a world-wide increase in the average level of education (Simon & Boggs, 1995; Goklany, 2000). As society becomes more wealthy and information becomes easier to distribute, educating people becomes both easier and more desirable. Illiteracy has virtually disappeared in the developed countries, and is strongly on the wane in the Third World. The average number of years that people spend in school has risen spectacularly, e.g. in the USA from about 4 in 1870 to about 18 in 1990 (Goklany, 2000). In the developed countries, a majority of pupils now goes on to higher education. The general consensus is that education will no longer be finished after college, but become a permanent process, as employees need constant training to keep up-to-date with the developments in their field and in society at large.

More important even than the quantity of education is its quality. Education in itself is of little avail if it is limited to the rote learning of religious books or ideological pamphlets. As uncountable observers have noted, the most critical capacity is learning how to learn. This requires education where the focus is not on static rules or facts, but on methods to autonomously analyse problems, find relevant information, synthesize the results, and thus develop new knowledge. An education for the 21st century will in particular need to teach us to better understand complex systems, and to avoid the typical errors of judgment that result from their counter-intuitive behavior (e.g. Casti, 1994). This requires an in-depth reflection on the observations, concepts and principles developed in disciplines such as systems theory, cybernetics, self-organization, chaos, evolution and complex adaptive systems (Heylighen et al., 1999).

While education is necessary to tackle the complexification of society, it is clearly not sufficient. Once we have reached the state of permanent education, further increases in the amount of education are no longer possible. Yet, it are precisely the most highly educated people, the managers, lecturers, scientists and technologists, that seem to suffer most acutely from information overload. While they may be objectively most competent to make decisions, they are subjectively most aware of the gap between the information that is out there and the one they can effectively assimilate.

The reason education must fall short of tackling the information explosion is biological: the human brain is an organ with a limited capacity for storing and processing information. Yet, this does not imply an old-fashioned biological determinism. Recent developments in psychology and neurophysiology paint a more optimistic picture: the brain is much more plastic than was thought, and there appear to be numerous ways to increase its capacity. This can be illustrated most concretely by examining *intelligence*.

While intelligence is a complex, multifactor phenomenon, which is undoubtedly affected by education, extensive statistical analysis of IQ test results points to at least one fundamental factor, the so-called *g-factor* (for “general” intelligence), that appears primarily biological (Jensen, 1998). The *g-factor* can perhaps best be understood as a measure of the efficiency of information processing in the brain. It is positively correlated with indicators such as the capacity of working memory, the speed of neuronal transmission, and the physical size of the brain (Jensen, 1998). While there is a clear genetic component to a person’s *g-factor* intelligence, there is also a strong influence from the environment beyond the effect of education.

This is shown most dramatically by the *Flynn effect*: average IQ scores for the population appear to have been rising with some 3 points per decade, and this for at least the past century (Flynn, 1987; Neisser, 1998). Surprisingly, this on-going increase is most pronounced for those—*g*-related—components of intelligence that are least dependent on school knowledge. This indicates a primarily biological phenomenon. Several factors have been proposed to explain this secular rise in intelligence (Neisser, 1998): richer nutrition, better health, more cognitive stimulation by a more information-rich environment, more parental attention invested in a smaller number of children, etc.

Perhaps the simplest way to understand how the brain can become more efficient is to compare it with other organs that consume a lot of energy: the muscles and the cardiovascular system. The general efficiency of these organs determines a person’s

physical fitness, that is, capacity for sustained physical effort. It is well-known that fitness depends on many factors: genes (some people are born athletes), food (you need to eat sufficient proteins to build muscle and sufficient calories for energy), general health (ill people have poor condition), training (the more you use a muscle, the stronger it becomes), development (people who have been training from an early age become more accomplished athletes), drugs (there exist plenty of legal and illegal performance-enhancing supplements), etc. Each of these factors known to increase physical fitness has an equivalent that may similarly increase mental fitness.

Most obviously, to develop and maintain an efficient brain, you need to eat richly and healthily and to train your “cognitive muscles” by constantly challenging them with new information and non-trivial problems. Evidence for the training effect can be found in a number of special programs to boost intelligence. While the children participating in such programs saw their IQ rise with several points compared to others, this advantage was lost within a year or so after the program had finished (Jensen, 1998). While this was interpreted as failure, it matches what we know about training for physical strength and endurance, whose effect also quickly dissipates after the training is stopped. As is well-known, both mental and physical capacities are subject to the rule of “use it or lose it”. Thus, it seems likely that the increasing supply of information and complex challenges that accompanies ephemeralization actually increases the mental fitness of the population. (It is ironic that at the same time ephemeralization has reduced the need for physical effort and therefore physical fitness.)

Another promising approach to boost intelligence are so-called “smart drugs” (Dean et al., 1991): pharmacological substances that apparently improve the efficiency of brain processes. Some of these substance have a very specific action in stimulating or inhibiting certain neuro-transmitters, and thus may perturb a finely tuned balance. Others, though, seem to have a more general “friction-reducing” effect: they facilitate blood circulation in the brain and thus the delivery of oxygen and nutrients to the highly energy-intensive neural processes; moreover, they typically neutralize the free radicals produced by this oxygen metabolism that would otherwise randomly perturb the ongoing processes and thus contribute to the background “noise” (which appears to correlate negatively with g-factor intelligence, Jensen, 1998). As a result, such antioxidant substances, like the bioflavonoid extracts from the Ginkgo tree, seem able to improve information processing and memory (Stough et al., 2001), and to combat fatigue and premature aging.

Just like athletic performances continue to improve year after year, there as yet does not seem to be a clear limit to the increase of “brain fitness” exemplified by the Flynn effect. Yet, this intelligence increase, while paralleling many of the physical processes of ephemeralization, does not seem strong enough to keep up with the information explosion. The reason is that IQ rises seem linear at best (with fixed increments), while the number of options that need to be considered increases exponentially. At best, the Flynn effect together with the rise in education levels may explain why information overload has not created any more serious problems yet in a society that has become immensely more complex over the past century. But it is clear that no further advances in neurophysiology or even genetic manipulation will be able to

boost intelligence beyond the physical limitations of the human brain, and those are obviously much stricter than any physical limitations on information networks.

3. The economy of attention

If we cannot sufficiently increase individual information-processing capacity, then perhaps we can boost the overall capacity of society by more efficiently prioritizing and allocating the tasks that need to be performed. The problem of information overload can also be formulated as *attention scarcity*: it is not so much that there is too much information, but too little time and mental energy to process it. The amount of cognitive effort or attention (Kahneman, 1973) that an individual can give to any issue is limited, and there are in general more issues that demand attention than attention that can be given. Therefore, attention is the true bottleneck, the one scarce resource on which all others depend, and thus the one that is intrinsically most valuable. While ephemeralization can amplify the availability of any other resource, it cannot augment the total amount of human attention.

Problems of scarcity are fundamentally addressed by *economics*, which is the study of how scarce resources can be allocated most efficiently. At the most basic level, “economy” means simply the careful management of resources, so that as little as possible is wasted. From this point of view, individuals should learn to optimally spend the limited amount of attention they have, by investing it only in the most worthwhile items. This means that the many messages or opportunities that are constantly clamoring for our attention should somehow be ordered as to their importance, so that the less relevant ones can be ignored (cf. Heylighen, 1994; Losee, 1989). While a reliable ordering would require paying detailed attention to all items in order to determine their relative importance, defeating the whole idea of saving on attention, there exist various rough-and-ready rules that can help us to estimate priority. All of us have to some degree developed an intuition that helps us to recognize what is important, but it seems worthwhile to try and formulate such rules more explicitly. For example, messages in the form of chain-letters or commercial announcements about products you do not plan to buy can be classified at the bottom of the priority list.

More fundamentally, the emerging science of *memetics* (Blackmore, 2000), which studies the spread of ideas, may provide a systematic list of criteria (Heylighen, 1998) that distinguish useful, informative messages from parasitic information or “mind viruses” (Brodie, 1996). The basic idea is that if two messages are equally common, they must have a similar overall score on the different selection or “fitness” criteria that measure the success rate of memes. However, if the first message scores high on the criteria that characterize successful “mind viruses” (e.g. it appeals to our innate fears and desires, it is impossible to falsify, it’s a “good story”, it demands people to further spread the message...), and the second one doesn’t, then it is more likely that the second one got there because it is truly informative, unlike the first one.

Distinguishing and filtering out unreliable or irrelevant information is one part of what Shenk (1997) calls *information hygiene*. The other part is avoiding to produce and transmit such information (cf. Sherwood, 2001). People should not only learn how to recognize information parasites and other forms of low-content messages, they should

themselves actively refrain from adding to this “data smog”. This implies among other things that people who want to communicate a message should learn to express themselves in a way that is clear, concise, to the point, and as much as possible based on reliable data, and that they should target their message only to the ones most likely to benefit from it. Only then will they minimally add to the existing information overload of their audience.

Agreeing about such “netiquette” or appropriate rules of conduct in communication may significantly reduce information pollution, but it will not stop the people who have something to gain in transmitting their messages. This is most obvious for commercial publicity, where it is in the interest of the seller to inform as many people as possible about their offers, but it also applies to individuals and organizations (e.g. lobbyists, pressure groups, political parties) who for various reasons want to attract attention to their ideas. While freedom of expression makes it impossible to strictly limit the number of messages that are produced, the concept of *attention economy* may suggest a more flexible approach. Ephemeralization has made the production and distribution of information extremely inexpensive, inciting senders to spread their messages ever more widely. It costs hardly anything to send a commercial message to millions (and soon billions) of email addresses. With such mass-mailings, while most addressees would find little of value in the message, only the tiniest response percentage is sufficient to make a huge profit. Therefore, there is no real advantage in targeting restricted groups. Where the cost for the sender is minimal, the cost for the receivers, while individually almost negligible, is collectively huge. Assume that an addressee spends on average a mere second to decide that a spam message should be deleted. If the message is sent to 100 million people, this entails a total loss of some 700 working weeks. Now consider the losses when 100 such messages are distributed every day!

The cost has shifted basically from sender to receiver. If attention is the most scarce, most precious resource to remain after ephemeralization, then it would seem logical that people should pay to receive it. While unorthodox, a straightforward way to implement this principle would be to instate an *information tax*. Instead of letting email be basically free, a protocol could be created so that every sender would pay a small amount (say, 10 dollar cent) per addressee. Such an amount would be too low to make anybody think twice about sending a message to a loved one, but it would make spamming uneconomical, forcing publicity messages to target their audience very precisely. The tax could be collected centrally, and used by the government e.g. for combating information overload at large. Alternatively, it could be implemented as a decentralized transaction, an “attention fee”, that is paid directly by the sender to the receiver. The protocol could be further expanded so that if the addressees of the message would indicate their satisfaction with the message (e.g. by clicking an “OK” button, or by maintaining a list of “OK” colleagues and friends), the fee would be waived. In that way, people would be less inclined to send messages that are unlikely to be appreciated, while the people who do get more messages than they desire would at least receive some form of monetary compensation for their wasted effort. (While the intention is different, there already exist schemes where people are being paid for their willingness to simply

pay attention to advertisements, e.g. by clicking on web banners, or listening to commercials during their phone conversations.)

This economic analysis of attention can be taken a step further. While attention is a universally valuable resource, some people's attention will be more valuable than others'. Generally, the attention of people who are powerful, popular or authoritative will be much more in demand, as their reaction to the messages they receive will generally have much more influence in the outside world. For that reason, presidents, film stars, religious leaders, royalty, and Nobel prize winners generally receive immensely more solicitations than little known pensioners or homeless people. According to the law of supply and demand, their attention should therefore command a much higher price. In practice, such people are surrounded by a secretarial staff that processes the great majority of the messages, and the upkeep of this staff does require a lot of money. The high status of these people is usually accompanied by material wealth sufficient to pay for such upkeep, and therefore there does not seem to be an urgent reason to force senders to pay huge sums in order for their messages to reach a high-status person. Moreover, such a purely monetary way of valuing scarce attention would appear highly undemocratic, making it almost impossible for non-wealthy people to get the attention of their leaders (though it must be noticed that in practice this is just what happens, even without explicit fees for attention-getting).

An additional argument why high-status people should not be paid more highly for their attention is that in a sense they are already being paid back by the attention they get themselves. Goldhaber (1997) has argued that attention is not only valuable because we have too little of it to give, but because it is intrinsically pleasant to receive. It is part of human psychology that we seek to increase our status, and this goes together with increasing the amount of attention we get from others. Therefore, becoming famous is the dream of many. Since ephemeralization has democratized wealth, but kept attention scarce, fame may actually have become more attractive than wealth. Goldhaber (1997) therefore suggests that the traditional economy, based on the exchange of material wealth, is being replaced by an economy based on the exchange of attention.

This view of the attention economy has a basic flaw, though: attention is not a tradeable good. While attention is valuable both when spending it and when receiving it, the one cannot compensate for the other. All the attention that is focused on a famous person's private and public life will not help that person tackling information overload. At best, public attention can be converted to money, as when it helps a pop star sell records, which in turn can help the receiver buy the support to process more information, but this seems hardly an efficient way to direct information processing capacity where it is most needed. The market's "invisible hand" that balances supply and demand may be a relatively effective mechanism for allocating tradeable goods and capital (cf. Heylighen, 1997), but the same does not seem to apply to attention.

One reason why attention is so difficult to allocate rationally is that people have very little control over the emotional drives, such as sex, status, and danger, that focus their attention on one subject rather than another. News and publicity agencies have very well learned how to manipulate these drives in order to sell their messages, e.g. by including pictures of sexy women or cute babies that are wholly irrelevant to the

message itself. Most of these drives are deeply rooted in our genes, being adapted to a prehistoric hunting-gathering lifestyle very different from our present information society. Yet, several authors (e.g. Stewart, 2000; Czikszentmihalyi, 1990), building on centuries-old spiritual traditions such as yoga, meditation and Zen Buddhism, have argued that it is both possible and desirable for people to learn to control these drives.

While the effort and discipline necessary to achieve mastery over one's emotions may be daunting, the first step is simply to become aware of the relatively simple ways in which our emotions are being manipulated. This awareness could be part of the rules of information hygiene that everybody should learn. Another reason why control over drives may not be so difficult to achieve is that, according to the need hierarchy of Maslow (1970, Heylighen, 1992), "lower", material needs become less pressing as they are better satisfied. Thus, in a society where most basic needs of food, security, company, etc. have been satisfied, people will spontaneously pay more attention to higher, cognitive needs. The problem remains that there is an "inertia of desire" (Heylighen & Bernheim, 2000) which keeps desires active long after the underlying needs have been satisfied. Here too, there may lie a role for a generalized education into "mental" hygiene.

4. Computational Intelligence

If individual information processing capacity is too limited, and the economic allocation of this capacity suffers from serious restraints, then we might try tackle the information explosion with the help of computer technology. It is clear that computers are much less limited in the sheer amount of information they can process, and that whatever limits there are now will vanish soon if we extrapolate Moore's law to the nearby future. While our conscious processing in short-term memory is extremely limited, it is clear that the more diffuse, automatic, subconscious processes relying on long-term memory (e.g. recognizing faces or producing speech) have a capacity that is still beyond the one of present-day computers. Yet, it is a fashionable exercise to chart the computing power of various organisms, from slugs to humans, together with the capacities of subsequent generations of computer processors. A simple extrapolation of the trend then invariably leads to the conclusion that human intelligence will be surpassed in a mere decade or two (see e.g. Kurzweil, 2000). The implication is, depending on the author's intrinsically optimistic or pessimistic world view, that either all problems will be solved for us by these superintelligent machines, or that these machines will take over command and relegate humanity to the "dustbin of history".

Though simple and attractive, this reasoning obscures a fundamental property of intelligence. While information processing, in the simple mechanical sense of manipulating bits or making computations, is a necessary component of intelligence, it is far from sufficient. No electronic calculator, however fast its circuits or extended its memory, can ever make a real-life decision. The domain of artificial intelligence (AI) has studied in depth how computers can be turned from mere calculators into devices with something approaching human intelligence. After many initial hopes got frustrated (e.g. automatic translation or general purpose reasoning), some hard lessons appear to have been learned. First, practical intelligence requires extensive knowledge; yet it is

extremely difficult to elicit this knowledge from human experts in a form sufficiently detailed and explicit to be programmed into a computer. This is the *knowledge acquisition bottleneck*. Second, true intelligence cannot be completely preprogrammed: it must be able to develop autonomously, to *self-organize*. Third, for an autonomously developing system to acquire knowledge about a realistically complex environment, it must be able to *interact* extensively with that environment, so that it can learn from its experiences (this requirement of interaction with a true environment is sometimes called “embodiment” or “situatedness”).

Such interaction requires very sophisticated sensors, that bring information into the system, and effectors, that execute actions so that the system can test its understanding of the outside world. In the realm of sensors (sensory organs and the perceptual apparatus that interprets their signals) and effectors (muscles, organs, limbs and the control systems that govern their action), human beings still far surpass any artificial system, as becomes immediately clear when watching the rigid, clumsy behavior of present-day robots. Moreover, even if it were possible to build robots with the same capacities for interaction as people, these robots would still have to undergo the variety of real-world experiences that people go through in order to build up a comparable level of knowledge and intuition. All in all, it does not seem worth the huge investment that would be needed in order to build a robot that at best can mimic the slow intellectual development of a child.

In conclusion, the successes of AI at present and for the foreseeable future seem restricted to specialized programs (e.g. expert systems, data mining, or chess-playing) that complement human intelligence in limited, well-defined domains. For real-world decision-making, on the other hand, human experience and intuition must enter the loop. The present drive therefore is not so much for independently intelligent programs, but for systems that support or “augment” human intelligence (IA, that is, Intelligence Amplification, rather than AI). The most popular paradigm at the moment is that of the software *agent*, a program that carries out various relatively simple tasks on behalf of its user, such as keeping track of contacts and appointments, seeking out items that the user is likely to appreciate, or negotiating with other agents, e.g. for buying an item at the lowest price. In spite of the hype, few of these promises have as yet been fulfilled. The reason is that because the agent’s intelligence is quite limited, it would only be able to solve such problems in a well-structured environment, where the different options and the relations between them would be unambiguously defined, using standardized codes that all agents could understand.

Creating such an environment is the main drive behind the vision of the *semantic web*, (Berners-Lee et al., 2001) a shared realm of networked information that is structured according to a consensual “ontology”, i.e. a taxonomy of object and relationship types, so that all users and their software agents would know precisely how to interpret otherwise ambiguous concepts, such as “address”, “profession”, “person”, etc. Yet, it seems that the effort needed to create a workable semantic web will be huge, and will need to be carried out mostly by humans rather than by computers. Moreover, however extensive an eventual formal ontology, there will always

remain an infinite realm of ambiguous, context-dependent situations that can only be tackled through real-world experience (cf. Heylighen, 1991).

A domain in which computer programs do seem to have made real contributions in tackling the information explosion is *information filtering*. The idea is that a program, “search engine”, or agent would help its user in selecting or prioritizing the available information by scanning all potentially relevant documents for those patterns or components that best match the user’s preferences. Processing megabytes of information to find a particular well-defined component (e.g. a keyword) is a task at which computers excel. While the human mind excels at getting the “big picture”, at developing an intuitive understanding of a complex situation based on a variety of fuzzy and subjective impressions, computers are best when working with huge amounts of discrete items, that must be classified, computed, or otherwise processed according to fixed rules. It is in this latter kind of task that the brain’s limits on memory and processing are most salient. Sieving through billions of items, and calculating the contribution of each of them to a given “preference function”, is typically a task that is impossible for the human brain, while being relatively easy for a computer. This explains the popularity of search engines, which are the preferred entry point for most people looking for information on the web.

Yet, the present search engine paradigm has intrinsic shortcomings, because of which it will never be able to solve the problem of information overload. Search engines rely on the presence of well-defined patterns or keywords to find the documents a user is looking for. This means that they necessarily fail when the specified keywords are absent, even though the document itself may be very relevant. Keywords may lack because the document does not contain any text (e.g. a picture, movie or sound), because the author of the document used different words to describe the subject (the problem of synonyms), or simply because the user does not know the exact keywords that would characterize the subject (s)he is interested in. While keyword search will therefore fail to retrieve many relevant documents (poor *recall*), it will moreover burden the user by retrieving plenty of irrelevant documents (poor *precision*). A document may contain the right keywords but be irrelevant because: 1) the keywords were used in a different sense (e.g. “kind” as “type” or “kind” as “friendly”, i.e. the problem of homonyms), 2) the author has repeatedly included popular keywords merely to increase the chances that the document would be retrieved, or 3) simply because the fact that a certain word is repeated many times in a text does not imply that that document contains information relevant to the problem at hand (e.g. a message that repeats “O God! O God! O God!” will have little relevance for theological questions). Truly grasping the relevance of a document requires not only the recognition of words, but the parsing of sentences, and most importantly the analysis of meaning, and this is precisely what cannot be done without a huge amount of real-world experience.

In conclusion, while keyword search and other methods based on the formal analysis of components can facilitate the selection of information, the final judge of the meaning, value or relevance of a document must be a human, using real-world experience and intuition.

5. An integrated approach

5.1. *Collective intelligence*

The previous section has proposed three basic components needed to tackle the information explosion: 1) individual human minds; 2) economical or social rules for the allocation of attention; 3) computer systems to support human decision-making. If you cannot solve a problem on your own, then perhaps you can delegate it to someone else (and if necessary pay that person for the invested effort), or have a computer program help you to sieve through the complexity. It was argued that each component on its own is insufficient to keep up with the information explosion. On the other hand, simply adding up their contributions will at most triple the overall power, again falling short of an exponential growth in complexity. We need to combine these components into a system that is *more than the sum of its parts*, in other words where the components work in *synergy*, multiplying each other's capacity.

There is a name for the synergetic use of individually intelligent components: *collective intelligence* (Lévy, 1997). The most famous examples are insect societies, such as ant nests, bee hives or termite colonies (Bonabeau et al., 1999), that consist of individually dumb components, but are capable of surprisingly smart behavior when working together. In an earlier paper (Heylighen, 1999), I have analysed some of the basic mechanisms underlying collective intelligence, and suggested how these mechanisms could be implemented to turn the world-wide web into an intelligent system. There is no space in the present paper to review the technical details of this analysis, but the basic issue is to efficiently coordinate the actions of many different components. The idea is that different individuals, agents or computer programs would contribute their specific knowledge, solve those partial problems or make those decisions for which they are most competent. The results of this cognitive effort would be shared with all other components in a coherent system that I have called a "collective mental map" (CMM). A CMM consists of cognitive resources (typically documents or database records, but this may also include computer programs, agents, and human experts), that are linked by a network of associations. This network would be organized in such a way as to minimize the effort in getting any resource to the place where it is needed.

A CMM for the whole of humanity would obviously be an enormously complex system. No system, human or technological, would be able to exert any form of centralized control over such a map so as to coordinate or allocate contributions. Any mechanism of coordination must be *distributed* over all contributing components. In other words, a CMM for global society must be *self-organizing*. Hints on how such a self-organizing mental map could function can be found both in the collective foraging behavior of ants, and in the organization of the brain. In both cases, paths (sequences of links) that lead to useful solutions are reinforced, while paths that lead to poor solutions are weakened and eventually erased. In both cases, if different paths lead to the same solution, the most efficient one is strengthened while the less efficient ones eventually lose out. In both cases, new paths or links are initially created locally, by individual ants

or between individual neurons, but if successful are amplified by a non-linear, positive feedback mechanism, so as to generate a potentially global order.

While ant foraging provides a concrete analogy to illustrate how high collective intelligence can emerge from much more limited individual intelligence, the CMM system we are trying to understand will be immensely more complex than any insect society. As the most complex system we know until now, the human brain provides a more realistic, albeit still limited, analogy. Therefore, my preferred metaphor for this encompassing intelligent system is the *global brain*. This term was apparently first used by Russell (1995), although many authors before and after have proposed related concepts, such as world brain, world mind, noosphere, super-brain, digital nervous system, super-organism etc. (Teilhard de Chardin, 1955; Stock, 1993; Mayer-Kress & Barczys, 1995; Heylighen & Bollen, 1996; de Rosnay, 2000; Goertzel, 2001; Heylighen, 2002). While most of the authors used these concepts in a purely metaphorical—or even metaphysical—sense, the quick development of digital network technologies and theoretical advances in domains such as complex systems, cognitive science and cybernetics allows us to be more concrete, and provide a first glimpse of how such a global brain might tackle the problems of complexity and information overload.

Another analogy helpful to understand such distributed control or self-organization is the market's *invisible hand*, which directs resources to the place where demand for them is highest (Heylighen, 1997). We have argued that a hypothetical market of attention lacks a form of stable tradeability necessary for the rational allocation of capabilities, since receiving attention is not in itself useful to meet one's demands. Attention only becomes valuable when it is used to make important decisions or solve significant problems. Similarly, in a traditional market, money does not have value in itself, but only when it is used to invest in worthwhile enterprises or commodities. The function of the market is not to trade money, but to facilitate exchanges of labor and goods, using money as a measure of value. Similarly, an effective attention economy would not so much *trade* attention, but *direct* attention to the issues where it can be applied most usefully, i.e. where "demand" for it is highest.

The problem is that attention, unlike money or other forms of capital, is not durable: whatever is not used the moment it becomes available, is thereby lost. You cannot save or accumulate attention for later. Therefore, directing attention to the most worthy issues must happen as quickly as possible if it is to be efficient. If you want to invest your scarce attention in reading a good book on tackling repetitive strain injury, but you first must wade through library catalogs, find and read various book reviews in magazines, or browse through the many sections and books in nearby bookshops, then you have squandered a valuable resource. A global brain-like network would immediately bring the most interesting document to your attention, highlighting the sections that are most relevant to your personal situation. Let us try to illustrate how such a system might function by examining a paradigmatic problem.

5.2. *Illustration: an intelligent transport system*

Ephemerization has made transport physically easier, but organizationally more precarious. Everywhere governments, organizations and individuals are struggling with problems of traffic congestion, pollution, delays, and noise. Public transport would appear to significantly reduce many of these problems, but is not very popular because of its rigidity. If you need to go someplace, you would rather be able to jump in your car and drive there the moment you are ready, instead of having to spend time planning connections, checking schedules, and ordering tickets, only to find out that there are no buses or trains running at that particular moment or to that particular destination. Let us try to imagine how a distributed information network could help tackle these issues.

An obvious first step is to make all public transport schedules (trains, buses, subway, ...) available on the net. Most railway and subway companies already have a website where you can enter the station from which you would depart, and the station nearest to your destination, and receive all the connections that would arrive around a particular time. This is still pretty rigid, since you may not know which stations are near to you or your destination, or how to get there. Moreover, you may need to combine different forms of transport, such as bus, subway and train, that have independent schedules.

Since Spring 2001, the Brussels company for public transport (www.mivb.be) has been offering a more intelligent support system. It allows you to enter your precise location and destination (in the form of a street address, or name of a landmark, such as “Museum of Modern Art”), without needing to worry about stations or bus stops. The system then calculates the quickest combined connection, taking into account the time schedules of all forms of public transport available (tram, bus, subway) and the time you would need to walk to and from the different stopping places. Moreover, it provides you with detailed guidelines, such as: “take the second side street on the left, ..., walk 100 meters to the station Delta, there take the subway at 11.47, get off at ..., ..., get on bus 47 at 12.03, ..., from the bus stop cross the street and walk to the right for 2 minutes, until you reach the museum at 12.29”. Since because of congestion buses or trams do not necessarily run as scheduled, the system includes a module that tracks their position in real time, and thus can warn you when they will arrive. The system finally allows you to specify your preference for the route that either is quickest overall, involves least walking, or has least stop-overs.

Such a system takes much of the drudgery out of planning a journey, and thus allows you to save on your valuable attention. However, by adding a few recently developed technologies, we could take it several steps further. First, the real-time capabilities become truly useful only when the system can be consulted in real time, i.e. while you are on the journey. This is already possible using a laptop or palmtop computer with a wireless internet connection. Moreover, such a device might include the kind of Global Positioning System that is already popular in cars, so that the system would know where you are, without need for you to enter an address. In this case, the guidelines could be constantly updated, so that if by error you walk down the wrong street, or another bus arrives than the one that was scheduled, the system can recompute the optimal route and guide you interactively to your destination.

Second, instead of you having to buy tickets, payment could be done automatically, by transferring the correct amount of digital cash from your account to the one of the transport company the moment you finish each leg of the journey. This would imply the additional option of choosing the least expensive route rather than the quickest. This becomes especially interesting if the system would include different companies or forms of transport (e.g. taxis vs. buses) with different prices and different transport offers. Depending on the route and the time of day, overall cost, comfort and duration would vary, but your personal agent would know your preferences, and negotiate with the transport system to find the option that best suits your situation and budget.

Until now, the system that we sketched optimizes travelling for the individual, thus minimizing the attention spent in planning and performing any journey. Yet, optimization can be extended to the collective level. First, by making public transport more attractive, the system would already reduce the pollution and congestion created by the less efficient private forms of transport. More fundamentally, an intelligent system could make public transport itself more productive. Imagine that all individuals would communicate with the transport system the way we sketched it. In that case the system would know precisely how many people travel between any two destinations at any time of the day. This would allow it to determine the most efficient way to organize the transport network. For example, if the system observes that many travellers' destination lies in between two stopping-places, it might decide to move a stopping place or create a new one. Or, if it notes that many journeys include subsequent segments of three different bus routes, it may create a direct bus route along that trajectory.

Such changes to existing schedules or routes are still rather rigid. A truly interactive system would moreover be able to adapt routes in real time, depending on the demand. A simple example can be found in the "group taxis" providing a flexible form of transport in many Third World countries. These are vans that can carry some dozen people and that drive around the more frequented roads picking up as many people as possible. The driver asks them for their destination, and if it is not too far out of the way from the one of the other passengers, he will take them in and adjust his route accordingly. The price is negotiated on the spot. This system is flexible thanks to the driver's experience, intuition and talent for improvisation.

An intelligent network should be able to provide the same flexibility, but with more dependability and efficiency. Roaming vans or buses would be directed by the system to the place where travellers are waiting, using the most efficient route that combines their various locations and destinations. Along busy stretches, the bus would be practically full, and each passenger would pay only a small share of the overall cost. Travellers who need to go to an isolated destination or who travel late at night, on the other hand, may find themselves alone with the driver. In that case, they may have to wait longer to be picked up, and pay a higher price, but still be assured that they will get the necessary transport. Again, price and waiting time can be negotiated between the traveller's agent and the system. In that way, the "invisible hand" of the market (albeit electronically supported) can adjust supply and demand of transportation in the most efficient way.

At this level of flexibility, the distinction between public and private transport becomes moot: such a system can also accommodate individual drivers, who are willing to share their car with passengers for part of their trip, in return for a fee negotiated between the respective agents. This is a more flexible version of car-pooling or hitchhiking. Since the overall intelligent transport system would know the identities and locations of driver and passenger (while protecting their privacy towards outsiders), this would moreover reduce the risks of crime associated with hitchhiking. Finally, the system could encompass car rental organizations, and direct the traveller to an empty vehicle for personal use rather than to a vehicle with driver, depending again on variables such as cost, location and personal preferences. The traveller's communication device would receive an electronic key from the network that would allow driving the car for as long as needed, in return for the negotiated transfer of electronic funds to the car's owners.

The proposed supply-and-demand driven merger of public and private transport does not imply that market forces should reign supreme. There are more benefits to public transport than the simple fact that carrying more passengers per vehicle reduces costs. Public transport can moreover reduce pollution, noise, energy consumption and congestion in a way that benefits society in the long term, without direct benefit to the traveller here and now. These collective benefits determine additional variables that should be taken into the equation when calculating the recommended route. They can be expressed as constraints that prohibit certain routes, e.g. a noisy bus passing near to a hospital, or a non-local car passing in front of a school entrance at the moment the children are coming out.

Others can be expressed as additional weights that bias the overall decision to the one that is optimal for society rather than for the individual traveller or transport firm. For example, the government can subsidize certain instances of transport and tax others, depending on their relative advantages and disadvantages ("externalities"). These corrections to the market price would automatically be taken into account when calculating the cost of the different options for the traveller. For example, the system may offer the train as cheapest option because it is less polluting, even though the immediate cost of a bus journey might be lower, or it may collect a toll on single passenger cars during peak hours.

Unlike traditional taxes and subsidies, such corrections should be able to adapt in real time, e.g. increasing the cost of more polluting options in proportion to the amount of air pollution that is registered by sensors. In that way, the present level of pollution could be regulated cybernetically, through a negative feedback-based control system (Heylighen, 1997). Moreover, the system would be able to react *immediately* to sudden perturbations, e.g. by redirecting traffic around the obstruction created by an accident. Thus, it would be able to prevent the self-reinforcing processes that lead to the build-up of a traffic jam, and that in a sufficiently dense traffic can be triggered by a mere local slowdown involving one or two cars.

Such an intelligent control of transport may seem very desirable, but computer scientists are likely to argue that optimizing such a system, while in principle possible, will in practice be far too complex for any existing or anticipated type of computer. It is

sufficient to note that the notorious “travelling salesman” problem, which is the standard example of the most difficult, “NP-complete” type of computation, is in fact a simplified version of the kind of optimization problem we are considering here. Happily, we do not need to find optimal solutions in order to achieve substantial improvements on the present system. The collective behavior of ants shows the way to a distributed, self-organizing approach that can achieve solutions near enough to the optimal one. Dorigo et al., (1996) have shown how “ant-based” algorithms are surprisingly effective in tackling the travelling salesman problem. The principle is that individual agents or “ants” heuristically try to find the best route for their leg of the overall journey, while their local contributions are superposed in a non-linear manner so as to generate a global trajectory that is shorter than any individual agent could have found. Projected back to the transport system, this means that the solution is not found by any central computer, but by the appropriate combination within a shared information system of the decisions made by all individual travellers and their software agents. The resulting overall trajectory is used to guide further travellers, who themselves contribute to the further improvement and updating of the trajectory, in a self-reinforcing loop.

5.3. *Optimizing production and services*

You may wonder why I have devoted so much space to discussing transport in a paper whose subject is the much more general issue of complexity and information overload in society. The reason is that transport proposes a concrete, easily visualizable illustration of the kind of problems I am trying to address. More importantly, the suggested solution to this paradigmatic example can be extended to the general problem.

The generalization from transport of people to transport of goods is straightforward: it suffices to replace the location, destination and preferences of the traveller by those of the supplier and client of the good. The generalization from transport to production requires a somewhat higher level of abstraction: every production process can be represented as a trajectory in *state space*, moving from the “raw” state of materials or components to the “finished” state via a number of intermediate states. This trajectory can be optimized taking into account the moment-to-moment supply of raw material and demand for finished goods. For example, when the electronic payment systems of booksellers register a higher demand for a particular novel, printing presses can be adjusted to immediately start producing more copies, while increasing the standing order for paper from the suppliers. Using principles of self-organization, such as ant algorithms and the law of supply and demand, optimizations for individual processes can be combined into a global optimization that takes into account mutual dependencies (e.g. the product of one process being used as a component for another process, or the same raw material being divided among different production processes). Effects at the collective level (e.g. some processes are more polluting or consume more scarce resources than others) can be taken into account by imposing additional constraints or preferences.

Services, such as medical treatment, too can be conceptualized as a trajectory from an initial state (e.g. the patient being ill) to a desired state (e.g. the patient being cured).

Again, an intelligent system can try to find the “shortest” route from the one state to the other, taking into account various constraints (e.g. avoiding risk to the patient’s life) and optimization criteria (e.g. minimizing cost, pain, and side-effects), at both the individual and collective level, while using the collective experience of other patients with similar ailments as stored in the shared database.

5.4. *Information routing*

Similar applications of these general principles can be found in the realm of knowledge and information. Most obvious is the routing of “raw” bits or signals through communication networks (Huitema, 2000). The TCP/IP protocol underlying the Internet already uses an elementary form of distributed self-organization, explaining the Internet’s remarkable flexibility and robustness. Ant-based algorithms are presently being investigated to optimize routing of calls through the telephone network so as avoid congestion (Schoonderwoerd et al., 1996).

More subtle is the minimization of information overload by the routing of messages or announcements to people: which addressees should receive which messages with which level of priority? In organizations, such as companies or administrations, where the different roles and responsibilities are well-defined, *workflow* systems may provide a partial solution (Schael, 1998). These are software programs following a predefined set of rules to ensure that the right request, task or announcement is sent to the right person at the right time. However, such systems tend to be rigid and not very intelligent in adapting to messages that do not fit the predefined patterns.

A more flexible and distributed approach is suggested by collaborative filtering (Shardanand & Maes, 1995): a person who has already read a message may score how relevant or interesting the message is for him or her. By correlating the scoring patterns of different people, the system can determine which people have similar interests. When new messages arrive, only a few people need to evaluate the message for the system to estimate how interesting the message is for which other people, thus establishing a priority ranking that allows filtering out the less informative messages. There are many refinements and extensions possible to this basic scheme (cf. Heylighen, 1999). The system could for example calculate the similarity of new messages to messages that have already been evaluated (taking into account variables such as sender, subject, and density of keywords) and thus estimate relevance before anybody has read the message. This would ensure that truly important messages (e.g. the director-general announcing a company reorganization) would immediately get noticed, while junk mail (e.g. an offer for free access to a porn website) would move to the bottom of the heap.

An important issue in this respect is the division of labor (cf. Heylighen, 1999): in what way can problem-solving or information-processing be divided most efficiently among the different individuals and non-human cognitive resources? This issue has recently become the object of empirical studies under the label of *distributed cognition* (Hollan, Hutchins & Kirsh, 2000).

5.5. *Collective Information retrieval*

After considering how *information* travels to people, we should also look at how *people* travel (“navigate”) toward information. In Internet parlance, the former is called “push” (information being sent to receivers), the latter “pull” (the users themselves requesting or retrieving information). This is the domain in which I have done most research (e.g. Heylighen, 1999, 2002; Bollen & Heylighen, 1998).

Search engine technologies which locate data on the basis of rigidly defined keywords must be supplemented by the collective intelligence of the users of the web. Again, the ant-trail paradigm is directly applicable: the more users have travelled a particular path, following hyperlinks or search results, from document A to document B, the stronger the mutual relevance of B to A can be estimated to be, and the stronger or more direct the link from A to B should become. I have called this approach “web learning”: the web learns new links between documents in the same way that the brain learns to create associations between phenomena experienced within a short time interval. Such learned links would create direct routes or shortcuts connecting documents that users frequently consult together. Complementarily, documents that are rarely used together would lose their direct links, thus minimizing the burden on the user who has to choose between a host of seemingly equivalent options.

Note that such a distributed reorganization of the information network is potentially much more powerful than a reorganization of the transport network, since the latter is subject to physical constraints that cannot be overcome by simple organizational measures. For example, no matter how efficient the transport system, you cannot create a 1 minute connection between two locations that are a hundred miles apart. On the other hand, two web pages that are a hundred links apart *can* be connected directly by a single link.

One advantage of this approach is that associations between documents do not need to be formally defined, like in the semantic web: it is sufficient that users intuitively establish some kind of association simply through the pattern of their usage. Thus, the system could create an associative link between a melancholy photo of a sunset and a melancholy jazz melody, without anybody needing to label either of them by a keyword or category such as “melancholy”. Therefore, users do not need to explicitly specify the information they are looking for: it is sufficient that they are able to recognize items related or similar to the ones they want to find for the system to guide them interactively to the desired pieces of information.

By using the technique of “spreading activation”, they do not even need to find any particular item strongly related to the thing they are looking for. They can start simply with mutually dissimilar items that all have some, possibly indirect, association with what they are looking for, as this is sufficient for the system to retrieve those documents that are most strongly related to all of them. For example, our experimental learning web system (Bollen & Heylighen, 1998) would retrieve the concept “office” when the user selects the concepts “building”, “work” and “paper”. With such systems, users do not even need to be able to formulate what they are looking for: it is sufficient that they can intuitively indicate whether something is more or less related to what they want. The applications are similar to those of collaborative filtering: the system may

recommend e.g. pieces of music or paintings based on the likes or dislikes of a user for other works of art, by inferring the implicit similarities between works of arts from the collective preference patterns of a large group of previous users (Shardanand & Maes, 1995; Heylighen, 1999).

If such a system were implemented on the level of the web as a whole, information overload would be largely eliminated, since all available options would constantly be prioritized as to their estimated degree of relevance for this particular user at this particular moment, without the users needing to do anything more than what they anyway do: looking at a piece of information and explicitly or implicitly indicating how interesting or relevant they find it. (Such implicit evaluation can e.g. be inferred simply from the time spent reading or using the document, Nichols, 1998; Claypool et al., 2000).

Such a system would optimize individual navigation paths through web space, similar to the way the previously sketched intelligent transport system would optimize journeys through physical space. Here too we could envisage adding optimization criteria at the collective—rather than individual—level, although there is generally much less need to constrain individuals in their use of information than to constrain them in their use of physical resources. “Hard” constraints would amount to a form of censorship, where the consultation of certain types of documents (e.g. child pornography, recipes for making a bomb, or manuals for terrorists) without special licence would be made impossible by the system.

“Soft” constraints would merely bias the optimization criteria so as to make certain types of information easier to retrieve, e.g. educational websites or government guidelines, and others more difficult to retrieve, e.g. tobacco advertisements, racist propaganda, or scientifically unfounded and potentially dangerous “cures” for various illnesses. This means that such documents would still be available for those who specifically want to consult them, but that users who just are browsing the web, following associative links, are unlikely to ever encounter them. Thus, naive users, such as children, would run a much smaller risk of being subjected to dangerously misleading information.

Obviously, whether such constraints should be imposed and, if so, what form they should take, will constitute a thorny political issue, demanding a deep and wide-ranging discussion. The semantic web-related efforts of the World-Web Consortium (W3C) already include a form of classification where certain categories of documents (e.g. pornography) could be defined as “out of bounds” for certain categories of users (e.g. children), but this approach suffers from the general difficulty of formal classification that I mentioned earlier, ignoring the fuzzy, subjective and context-dependent character of categories such as “pornography”.

Adding outside optimization criteria to the self-organization of links in the web, thus reinforcing or reducing the strength of certain links, seems to provide a gentler and more flexible option. This could be done e.g. by downgrading the weight of links to documents that contain certain keywords (e.g. offensive words or different labels for tobacco and illegal drugs), or, more reliably, by allowing a certain category of “responsible” users (e.g. parents, experts, or a government-approved board of

evaluators) to more strongly influence the weight of certain links than the average public would. In that way, differing opinions on the dangers of a particular document could balance each other out, and document weights could vary depending on the context from which one approaches them. E.g. a quack cure for cancer would be virtually unreachable from a page listing what to do in case you are diagnosed with cancer, but might have a direct link from a page on crackpot schemes and pseudo-science.

5.6. *The information market*

It is worth noting that soft constraints already exist in that search engines change the position of websites in the list of search results depending on the amount of money the website's owner is willing to pay. Such stealth manipulations of information retrieval are worrying, as they abandon relevance or value of information for purely commercial interests, where it is the highest bidder who determines what information is most likely to be read by the public. Tackling these will require some form of legal control at the level of society, where the criteria for optimization are open for discussion by the public.

A much fairer and more transparent way to reconcile commercial interests with constraints on information retrieval is the instauration of *micro-payments* for document consultation (see e.g. Nielsen, 1998). A general complaint about the web is that it is very difficult to make profits in that medium. This in part explains the large number of bankruptcies among "dotcom" firms following the Internet boom. Most of these firms started from the idea of making money through *advertisement*, where one website (e.g. a free online magazine) would be paid for the link it provides to another website (e.g. a bookseller). The problem is that with an overload of documents and links on the web, users are not inclined to click on an ad simply because it appears on a page they find interesting. This led to very disappointing "click-through" statistics, and the unwillingness of advertisers to pay much money for such ads. As a result, many free websites could not earn enough money to pay the costs for gathering and editing the information they provide. The alternative approach, websites for which you have to register and pay, get a disappointing number of registrations, simply because users don't find it worth the effort to make a complicated and costly transaction in order to get access to information of which they only need a small part and which they may well be able to find for free elsewhere.

The solution is to let users pay "per view", that is per page consulted, but this in a totally transparent way without requiring any more effort than clicking on an ordinary, non-paying link. This can be achieved by loading the user's browser with a certain amount of digital cash, a small amount of which is automatically transferred to the owner of a website each time a document from that website is downloaded. The user—or the user's agent—would be aware of the cost of the consultation, e.g. through a bar in the browser window, the color or length of which would indicate the relative cost for each link that the mouse passes over. Thus, a user or agent could decide that a particular link is too expensive with respect to its *expected utility*, and therefore select another link.

An even fairer scheme—albeit more complicated to implement—would be to let users decide how much they pay *after* they have consulted a document, so that only

effectiveutility would be rewarded. To avoid users behaving as if nothing they read is interesting enough to deserve payment, users would be billed a default amount for every page they consult, but would afterwards be able to redistribute this already paid money by rewarding or punishing unexpectedly good or bad pages. This would make it difficult to earn money from misleading pointers.

Such transparent payment methods would make the law of supply and demand apply fully to the information available on the web. Thus, the web would profit from the power for self-organization provided by the invisible hand, where providers would be stimulated to supply the information that is most in demand, and where poor quality, uninteresting websites would automatically lose market share. Competition would moreover force information providers to adjust their prices downward. Since a large share of the information, supplied by the public and by not-for-profit organizations such as universities, government agencies, and professional associations, would remain available free (though voluntary donations could be accepted), this means that the average price of a document would be too small (e.g. a couple of dollar or euro cent) to make any user think twice about downloading it (Nielsen, 1998). Yet, large websites with high quality documents that get millions of downloads per week could still make enough money so that they could pay experts to gather and edit their information. In such an information market, like in the intelligent transport system, users could navigate automatically, with minimal conscious decision-making, by relying on agents that know their preferences about content and cost.

6. Conclusion

This paper has considered the most fundamental ways to tackle the problems caused by information overload and complexity. Increasing capacity by augmenting individual knowledge and intelligence is the most straightforward approach, but cannot be sufficient because of the intrinsic limitations of the human brain. Collective capacity can be increased by more efficiently allocating decision-making among individuals. This may be achieved by developing rules of information hygiene and an economy of attention. Information processing capacity can be further augmented by complementing human decision-making with computer support. However, the hard lessons from AI have taught us that computers alone cannot make important real-world decisions, and that human attention must remain in the loop.

The solution proposed in this paper is the integration of the three basic resources: human intelligence, computer intelligence, and coordination mechanisms that direct an issue to the cognitive resource (document, person, or computer program) most fit to address it. This requires a distributed, self-organizing system, formed by all individuals, computers and the communication links that connect them. The self-organization can be achieved by algorithms similar to those underlying the learning of associations in the brain, the laying of trails by ants, or the invisible hand of the market. The effect is to superpose the contributions of many different human and computer agents into a collective “mental map” that links all cognitive and physical resources in the most efficient way possible.

The resulting information system would be available always and everywhere, reacting immediately to any request for guidance or any change in the situation. It would constantly be fed with new information, from its myriad human users and computer agents, which it would take into account to find the best possible ways to achieve any task it is confronted with. Optimization would take place both at the level of the individual who makes the request, and at the level of society which tries to minimize the conflicts between the desires of its different members and to aim at long term, global progress while as much as possible protecting individual freedom and privacy. Such an intelligent, adaptive, “omniscient” system can perhaps be best understood through the metaphor of the *global brain*.

Such a global brain would solve the problem of information overload at the most fundamental level. For the individual it would provide constant decision-support, presenting recommended choices in their apparent order of importance, as determined by the individual’s own preferences, the experience of all other agents, and the collective preferences of society. Complemented by an in-depth teaching and implementation of the rules of information hygiene, such a system should be able to eliminate the stress of not being able to cope with the number of available options, while minimizing the risk of bad decisions because of insufficient information being taken into account.

Depending on the amount of attention or effort the individual is willing to invest in the decision, (s)he could either immediately accept the “default” recommendation, or examine a variable number of options in more depth, making a choice that is potentially very different from the one initially recommended by the system. For example, if you need to decide which bus to take to quickly get to a meeting, you probably will gladly follow the advice of the system, without further thought. But if you are planning to buy a house, you would rather spend a lot of time formulating your preferences and constraints, collecting a variety of offers, and visiting the most promising of those in person. The degree to which individuals actively participate in the decision-making will moreover depend on their own level of education and intelligence. The smarter and the more experienced in the domain you are, the higher your chances to find a solution that is better than the default one, and the more the system can learn from your contribution. Thus, individual advances in knowledge and intelligence will directly or indirectly benefit the capacities of the collective.

On the level of society, a global brain-like system should be able to tackle the problems associated with the complexity and sensitivity of causal networks. It should be able to take into account myriad events and their interactions, and intervene immediately, before any problem can snowball out of control. Of course, even an intelligence of this supra-human order of magnitude will never be able to reliably predict the future behavior of a system that is in essence chaotic. Yet, as cybernetics has taught us, detailed prediction (feedforward) is not necessary as long as the regulatory system can react quickly and adequately to any potentially dangerous perturbation (feedback) (Heylighen, 1997). Moreover, by modelling the effects of complex interactions such a system should be able to solve the productivity paradox, overcoming the socio-institutional inertia, the non-linear side-effects and bottlenecks that hold back productivity growth (Heylighen, prev. paper). This should allow the economy to grow

at a much faster and more stable pace. By optimizing at the global or collective level, this would moreover help us to eliminate inequality, poverty and underdevelopment, while eliminating most pollution, scarcities and waste.

Formulated at this level of abstraction, tackling all of society's problems seems almost too easy. In practice, the global brain envisaged here is unimaginably complex, and is likely to evolve slowly and haphazardly, requiring plenty of trial-and-error, and decades of hard work by millions of highly educated and committed individuals. Yet, it is my conviction that the major pieces of the puzzle—the theories and technologies necessary to design and build such a system—already exist. What remains to be done is to put them together. No individual, organization or government can plan, control, or supervise such an immense task: the jigsaw will have to assemble spontaneously. While such self-organization may seem like wishful thinking, several authors (e.g. Stewart, 2001; Wright, 2000; Heylighen, 2002; de Rosnay, 2001) have argued for the existence of evolutionary mechanisms that drive the development of such a global, cooperative intelligence. Only time can tell in how far these models will turn out to be realistic.

References

- Berners-Lee, T., Hendler, J., & Lassila, O. (2001): The Semantic Web. *Scientific American*, 284(5): p. 35-43 <<http://www.scientificamerican.com/2001/0501issue/0501berners-lee.html>>
- Blackmore S. (2000): *The Meme Machine*, (Oxford Univ. Press).
- Bollen, J. & Heylighen, F. (1998): A system to restructure hypertext networks into valid user models, *New Review of HyperMedia and Multimedia*, 189-213.
- Bonabeau E., Dorigo M. & Theraulaz G. (1999): *Swarm Intelligence: From Natural to Artificial Systems* (Oxford Univ Press).
- Brodie R. (1996): *Virus of the Mind: the new science of the meme*. Seattle, Integral Press, .
- Casti J.L. (1994): *Complexification: explaining a paradoxical world through the science of surprise*, HarperCollins, .
- Claypool M., Phong L., Waseda M. & Brown D. (2000): Implicit Interest indicators (WPI-CS-TR-00-18, technical report Computer Science Dep., Worcester Polytechnic).
- Csikszentmihalyi M. (1990) *Flow: The Psychology of Optimal Experience*. New York: Harper and Row
- de Rosnay, J. (2000) *The Symbiotic Man*. McGraw-Hill.
- Dean W., Morgenthaler J. & Fowkes S. W. (1991): *Smart Drugs and Nutrients* (Smart Publications).
- Dorigo M., Maniezzo V. and Colomi A. (1996), The Ant System: optimization by a colony of cooperating agents, *IEEE Transactions on Systems, Man and Cybernetics-Part B*, 26(1), p. 1-13.
- Flynn J.R. (1987): "Massive IQ gains in 14 nations", *Psychological Bulletin* 101:2, p. 171-191.
- Goertzel, B. (2001) *Creating Internet Intelligence: Wild Computing, Distributed Digital Consciousness, and the Emerging Global Brain* (Plenum)
- Goklany, I. M. (2000): "Economic Growth and the State of Humanity", PERC Policy Series <<http://www.perc.org/ps21.pdf>>.
- Goldhaber M. (1997): The Attention Economy and the Net. *First Monday* 2, No 4, <http://www.firstmonday.dk/>

- Heylighen F. & Bernheim J. (2000): "Global Progress II: evolutionary mechanisms and their side-effects", *Journal of Happiness Studies* 1(3), p. 351-374
- Heylighen F. & Bollen J. (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), p. 917-922.
- Heylighen F., J. Bollen & A. Riegler (eds.) (1999): "The Evolution of Complexity" (Kluwer, Dordrecht).
- Heylighen F. (1991): "Design of a Hypermedia Interface Translating between Associative and Formal Representations", *International Journal of Man-Machine Studies* 35, 491-515.
- Heylighen F. (1992): "A Cognitive-Systemic Reconstruction of Maslow's Theory of Self-Actualization", *Behavioral Science* 37, p. 39-58.
- Heylighen F. (1994) "Fitness as Default: the evolutionary basis for cognitive complexity reduction", in: *Cybernetics and Systems '94*, R. Trappl (ed.), (World Science, Singapore), p.1595-1602.
- Heylighen F. (1997): "The Economy as a Distributed, Learning Control System", *Communication & Cognition- AI* 13, nos. 2-3, p. 207-224.
- Heylighen F. (1998): "What makes a meme successful?", in: *Proc. 16th Int. Congress on Cybernetics* (Association Internat. de Cybernétique, Namur), p. 423-418.
- Heylighen F. (1999) "Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map", *Computational and Mathematical Theory of Organizations* 5(3), 253-280.
- Heylighen F. (2001): "Bootstrapping knowledge representations: from entailment meshes via semantic nets to learning webs", *Kybernetes* 30 (5/6), p. 691-722.
- Heylighen F. (2002): "The Global Superorganism: an evolutionary-cybernetic model of the emerging network society", *Journal of Social and Evolutionary Systems*
- Heylighen F. (submitted): "Complexity and Information Overload in Society: why increasing efficiency leads to decreasing control"
- Hollan J., Hutchins E. & Kirsh D. (2000): Distributed Cognition: toward a new foundation for human-computer interaction research, *ACM Transactions on Human-Computer Interaction* 7(2), p.174-196.
- Huitema C. (2000): *Routing in the Internet*, Prentice Hall
- Jensen, A. R. (1998). *The g factor: The science of mental ability*, (Praeger, Westport, Connecticut)
- Kahneman (1973) *Attention and Effort* (Prentice-Hall, Englewood Cliffs NJ)
- Kurzweil R. (2000) *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, Penguin
- Lévy P. (1997), *Collective Intelligence: Mankind's Emerging World in Cyberspace* (Plenum, New York)
- Loosee (1989): Minimizing Information Overload: the ranking of electronic messages, *Journal of Information Science* 15 (3), p. 179-189.
- Maslow A. (1970): *Motivation and Personality* (2nd ed.), (Harper & Row, New York).
- Mayer-Kress, G. & Barczys C. (1995) "The Global Brain as an Emergent Structure from the Worldwide Computing Network, and its Implications for Modelling", *The Information Society* 11:1, 1-28.
- Neisser U. (ed.) (1998): *The Rising Curve: Long-Term Gains in IQ and Related Measures* (American Psychological Association).
- Nichols, D.M. (1998) Implicit Rating and Filtering, *Proc. Fifth DELOS Workshop on Filtering and Collaborative Filtering*, Budapest, Hungary, 10-12 November 1997, ERCIM, 31-36.
- Nielsen J. (1998): "The Case for Micropayments," *Jacob Nielsen's Alertbox*, Jan. 25., <<http://www.useit.com/alertbox/980125.html>>.
- Russell, P. (1995) *The Global Brain Awakens*., Miles River Press.

- Schael T. (1998): *Workflow Management Systems for Process Organisations*, Springer Verlag.
- Schoonderwoerd, R., Holland, O.E., Bruten, J.L., & Rothkrantz, L.J.M. (1996) Ant-based load balancing in telecommunications networks, *Adaptive Behavior*, 5(2), 169-207.
- Shardanand, U. & Maes P. (1995) "Social information filtering", *Proceedings of CHI'95 -- Human Factors in Computing Systems*, 210-217.
- Shenk D. (1997): *Data Smog: Surviving the Information Glut* (Harper, San Francisco).
- Sherwood, K. D. (2001): *Overcome Email Overload with Eudora 5*, (World Wide Webfoot Press Palo Alto, CA).
- Simon J. L. & Boggs R. (1995): "Trends in the quantities of education—USA and elsewhere", in: Simon J. L. (ed.), p. 224-230.
- Stewart, J. E. (2000): *Evolution's Arrow: The direction of evolution and the future of humanity* (Chapman Press, Australia) <<http://www4.tpg.com.au/users/jes999/>>
- Stock, G. (1993) *Metaman: the merging of humans and machines into a global superorganism*. Simon & Schuster, New York.
- Stough C., Clarke J., Lloyd J. & Nathan PJ. (2001): Neuropsychological changes after 30-day Ginkgo biloba administration in healthy participants. *International Journal of Neuropsychopharmacology* 4(2), p. 131-134.
- Teilhard de Chardin, P. (1955) *Le Phénomène Humain*. Seuil, Paris. (translated as: *The Phenomenon of Man* (1959). Harper & Row, New York).
- Wright, R. (2000). *Non-Zero. The Logic of Human Destiny* (Pantheon Books).