PERGAMON

Engineering Applications of Artificial Intelligence 12 (1999) 389–399

Engineering Applications of
ARTIFICIAL INTELLIGENCE

Mind and autonomy in engineered biosystems

O.G. Clark[a,*], R. Kok[a], R. Lacroix[b]

[a]Department of Agricultural and Biosystems Engineering, Macdonald Campus of McGill University, 21 111 Lakeshore Road, Ste-Anne-de-Bellevue, Québec, Canada, H9X 3V9
[b]Department of Animal Science, Macdonald Campus of McGill University, 21 111 Lakeshore Road, Ste-Anne-de-Bellevue, Québec, Canada, H9X 3V9

Received 1 December 1997; accepted 1 January 1999

Abstract

Biosystems are unitary entities that are alive to some degree as a system. They occur at scales ranging from the molecular to the biospheric, and can be of natural, artificial or combined origin. The engineering of biosystems involves one or more of the activities of design, construction, operation, maintenance, repair, and upgrading. Engineering is usually done in order to achieve certain preconceived objectives by ensuring that the resultant systems possess particular features. This article concerns the engineering of biosystems so that they will be somewhat autonomous, or able to pursue their own goals in a dynamic environment. Central themes include: the computational abilities of a system; the virtual machinery, such as algorithms, that underlie these abilities (mind); and the actual computation that is performed (mentation). A significantly autonomous biosystem must be engineered to possess particular sets of computational abilities (faculties). These must be of sufficient sophistication (intelligence) to support the maintenance and use of a self-referencing internal model (consciousness), thereby increasing the potential for autonomy. Examples refer primarily to engineered ecosystems combined with technological control networks (ecocyborgs). The discussion is focused on clear working definitions of these concepts, and their integration into a coherent lexicon, which has been lacking until now, and the exposition of an accompanying philosophy that is relevant to the engineering of the virtual aspects of biosystems.

Keywords: Mind; Intelligence; Consciousness; Autonomy; Biosystems; Ecocyborgs; Engineering; Philosophy

1. Introduction

This paper comprises a philosophical and lexical basis for engineering the minds of highly autonomous biosystems. Biosystems are collections of physical and virtual components that perform together as integrated, living units. They range in organizational scale from the molecular to the biospheric and, as well, vary greatly in their degree of autonomy. The discussion presented here is general, but is illustrated with reference to a particular kind of biosystem called an *ecocyborg* (Parrott et al., 1996). Ecocyborgs consist of both biological and technological components that interact at the scale of an ecosystem, where the latter is defined as a community of organisms, together with their abiotic surroundings. Biosystems of this type can be engineered for a variety of purposes, which may be best served by tailoring their computational abilities; i.e., their capacity to transform input signals from their surroundings into output signals.

Currently, ecocyborgs are usually artificial in origin, or are derived from natural ecosystems by human modification. Humans have historically modified ecosystems to favor their own survival, and this has in part allowed them to expand their range outside of the ancestral environment to which they are evolutionarily adapted. They have accomplished this by introducing and extirpating species, supplementing soil nutrients, and altering the hydrological properties of watersheds,

* Corresponding author. Tel.: +1-514-398-7590; fax: +1-514-398-8387.
E-mail address: Clark@agreng.lan.mcgill.ca (O.G. Clark)

0952-1976/99/$ - see front matter
PII: S0952-1976(99)00010-X

for instance. Such activities form the basis of modern industries like agriculture, aquaculture, and silviculture. Insofar as these practices involve the modification of biosystems in pursuit of particular objectives, they can be considered as examples of biosystem engineering at the ecosystem scale.

The modification of ecosystems, as it has been practiced, is a primary reason for the rapid growth of the human population that has occurred during recent millennia. Human activities are, in turn, altering the Earth's ecosystems more rapidly and on a larger scale than ever before. The magnitude of these alterations is such that ignorance or carelessness could potentially affect the integrity of the biosphere. The changes that could result from further human activity should therefore be carefully considered, as should the ongoing impact of changes that have already taken place. The modification of ecosystems on such vast scales must proceed with attention to design, construction, operation, maintenance, and repair considerations, aspects of engineering practice that have until now been largely neglected when dealing with biosystems of this class.

In the short term, biosystems engineering principles could be used to moderate environmental crises by making ecosystems more *autonomous*, or independent in the establishment and pursuit of their own goals. This would increase their persistence in the face of external perturbations or the self-serving activities of component species such as humans. The idea of engineering ecosystems in this way is new, and until now has not been framed in the context of biosystems. It implies modifying their computational abilities, or the manner in which the pattern of interactions between their components transforms input signals into outputs. All biosystems have some ability for computation, but natural ecosystems are incapable of the abstract mentation necessary for significant autonomy. Ecosystems might, however, be endowed with the required abilities by transforming them into ecocyborgs through the addition of technological components. An ecocyborg could be engineered to have computational abilities of the appropriate type and sophistication for *consciousness*, meaning that it would be aware of itself to some degree in the context of its environment. This in turn would increase its potential for autonomy. This approach, and the lexicon that is developed here, could prove to be valuable in the engineering and sustainable management of Earth's ecosystems.

In the long term, the engineering of biosystems at the ecosystem scale not only could help to safeguard against environmental crises, but might also provide for the continued growth and survival of the human species. Expansion into space, for example, will be necessary if humanity is to continue to increase, simply because the Earth's finite resources cannot sustain perpetual growth. Moreover, planet-bound life is vulnerable in the face of planetary events such as collisions between asteroids and the Earth (Sagan, 1994). The establishment of self-sustaining colonies in space would provide practically limitless room for growth, and would better ensure the security of the species. Since people can only exist in an appropriate environment, extraterrestrial expansion will require the creation of artificial ecosystems that include humans. These will undoubtedly include many technological components, making them ecocyborgs. Moreover, since they will have to be self-sustaining in the isolation of space, they will have to be engineered to be highly autonomous. The survival of space-borne ecosystems would be more secure if their autonomy were independent of humans, since the ecocyborgs would then be able to function even if human guidance became impossible or ineffective. This might occur if the occupants were incapacitated or neglectful; it is also entirely possible that such ecocyborgs would simply be too large and complicated to be effectively controlled entirely by humans. The International Space Station (ISS), of which construction began in 1998, is an example of such a space-bound ecocyborg. The philosophy and lexicon presented in this article could be useful conceptual aids to engineering the ISS and its successors as viable, integrated, goal-oriented biosystems that include humans as components.

The cyborging of ecosystems illustrates how one class of biosystems might be engineered to be highly autonomous, but many of the concepts related to such an exercise are also relevant to the engineering of a much broader class of biosystems. Until now, these concepts have not been clearly defined as part of a coherent and useful lexicon such as the one presented here. In this paper, each concept is first discussed in broad terms, and then illustrated with examples. Frequent reference is made to animals, especially humans, since they are the most accessible and intensively studied autonomous entities. The ideas are then expounded in the specific context of ecocyborgs, and integrated into a conceptual framework that facilitates the engineering of these and other kinds of significantly autonomous systems. Although the framework presented here is loosely based on traditional human psychology, it is certainly not the only approach that might be appropriate. Since large-scale biosystems such as ecocyborgs are often composed of semi-independent agents, a serviceable framework might also be developed, for instance, from a sociological perspective.

2. Implementing mind in biosystems

In this paper, the *mind* of a biosystem is defined as the virtual machinery, including algorithms, that make possible all of its computational abilities. All biosystems possess some computational abilities, but these abilities, the virtual machinery that gives rise to them, and the physical substrates in which that machinery is embodied, can all differ greatly from one biosystem to another. Humans, for instance, possess a nervous system comprising specialized organs that embody highly adapted virtual machinery. This machinery gives rise to very specific computational abilities that make possible some degree of consciousness and autonomy. Natural ecosystems do not have such specialized structures, and so do not possess the kind of minds that humans do. Instead, their computational abilities reside in the way that input signals are transformed into outputs through processes such as interactions between the populations of their constituent species, the cycling of nutrients, and subtle phenomena like the transport of biologically active trace chemicals (McNaughton and Coughenour, 1981; Patten and Odum, 1981). The cumulative result of these processes is certainly computationally complex, but it does not make natural ecosystems conscious or autonomous in the sense that a human is. The virtual machines corresponding to these processes are more analogous to those embodied in the workings of the human digestive and circulatory systems than to those of the nervous system. Thus, this virtual machinery all gives rise to computational abilities, but is not generally considered to contribute to the capability of natural ecosystems to model or reason about themselves in the context of their surroundings (Engelberg and Boyarsky, 1979). They cannot, therefore, establish and work towards their own goals. They can, however, serve as a basis for engineered biosystems of greater consciousness and autonomy.

Biosystems can be engineered to have minds similar to those of humans. Ecosystems, for instance, might be endowed with an infrastructure to support the computational abilities required for high degrees of consciousness and autonomy. This can be done by including components that are not native to a naturally occurring ecosystem. The resulting comportment is then a consequence of both the inherent dynamics of the natural ecosystem segment and the influence of the additional computational components. If the latter are added to a biosystem with the express intent of regulating its comportment so as to achieve particular goals, then the exercise is one of guidance or control. Control can be intrinsic or extrinsic, depending on the conceptual boundaries that are defined. If the guidance components are considered to be internal to the biosystem, then the control is *intrinsic*, whereas if they are considered to exercise a controlling influence from outside the system boundary, then the control is *extrinsic*. Components called *perceptors* sense signals in their surroundings, and create information corresponding to values of the *observed* variables. *Control mechanisms* structure this information and devise strategies to keep the values of certain *controlled* variables within a particular range. Lastly, *effectors* implement these strategies by parsing them into the values of *manipulated* variables, or directives that induce final control elements to generate output signals (Kok and Lacroix, 1993). In expansive systems these components are often arranged in distributed networks, being widely separated in space but still linked together by communications channels so as to influence each other's activities (Kok and Lacroix, 1993).

The control of large-scale biosystems can be illustrated with the example of human intervention in ecosystems. As discussed previously, humans habitually exercise control over ecosystems in order to improve their own welfare in the short term. It is believed that modern humans are relative newcomers to most parts of the world, having spread from the African continent only during the last two-hundred-thousand years or less (Vigilant et al., 1991). They have inserted themselves into a variety of ecosystems to which they were not originally native, and now regulate these in order to meet their own needs. Humans thereby guide the ecosystems by acting as perceptors, control mechanisms, and effectors. Whether this guidance is considered to be either intrinsic or extrinsic depends on whether or not the humans are included in the ecosystem definition.

Ecosystems might also be engineered to be highly autonomous by cyborging them with technological devices. Computer control systems already endow some greenhouses and industrial fermentation units, for instance, with a modest degree of autonomy. This approach might also be applied to other ecosystems, with machinery replacing human muscles as effectors, electronic instrumentation performing sensory tasks, and computers acting as control mechanisms. These technological components would endow the ecosystems with minds that would increase the independence of their comportment, enabling them to guide themselves toward particular goals. In the future, cyborged ecosystems might serve as habitats in the human colonization of space, and the entire biosphere of the Earth might someday gain greater autonomy through cyborging with sensory and communications networks (Dyson, 1997).

Two general examples have been given of how the minds of biosystems might be engineered by including some kind of control system. Many other methods might be described that apply to different kinds of biosystems, and the advent of new technologies and novel applications in the future will make possible the creation of biosystems that cannot be foreseen today. It is therefore important to be able to discuss the concepts associated with mind in biosystems in a manner that is relatively context-independent. The adoption of several

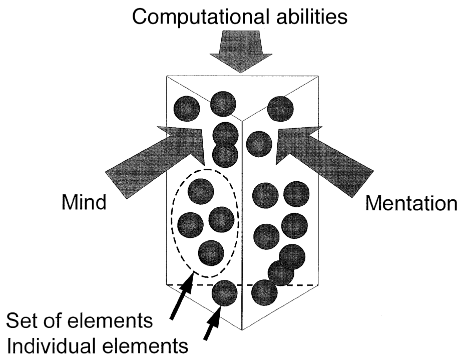

Fig. 1. Three-fold perspective of a computational entity.

complementary perspectives on the computational essence of mind can facilitate this objectivity.

3. Perspectives on computation

Three alternative perspectives are presented here that can be employed when discussing computation in biosystems (Fig. 1). The first refers to the *virtual machinery* that endows a biosystem with its computational abilities. This is the *mind* of the biosystem, and includes formal methods such as algorithms, although it may not be limited to these. The mind is referred to as being *virtual* because it creates, communicates, and manipulates information, but it must, nevertheless, reside on a physical substrate such as a brain or a landscape. This substrate might affect the performance of the mind, but in theory this is of only incidental importance. For example, an abacus, a Babbage engine and an electronic calculator can all potentially be used to perform the same mathematical operations. Although the speed of the operation might vary according to the instrument, the formal methods that are used can be qualitatively the same in each case. In an ecological context, for example, a rainfall event in a watershed might be transformed into discharge. The transfer function that relates the rainfall and the discharge might be the same for two different watersheds, but be mediated by different physical structures. In a completely natural ecosystem the transfer function might depend on topology and the hydrological characteristics of the soil, whereas in an ecocyborg it might result from the actions of a computerized network of drainage canals and hydraulic control structures that might be considered as extrinsic to the ecosystem. What is the same in both cases is the formal method, or virtual machine, that generates the output from the input.

Virtual machinery can be grouped into sets and supersets on the basis of the functions to which they give rise. Instructions, for example, are the most basic embodiment of virtual machinery in the context of the digital computers that constitute the computational infrastructure of current ecocyborgs (e.g., automated greenhouses). Instructions can be grouped into subroutines, and the subroutines into programs that can perform particular tasks. The boundaries of these sets and supersets are, however, arbitrary and can overlap. The same subroutine, for instance, might be used in several different programs. The virtual machinery of future ecocyborgs might be organized less like the linearly structured program code that is currently common, and more like natural biological mechanisms. As a case in point, artificial neural networks already exist that are modeled after biological nervous systems. As well, evolutionary programming techniques have been developed, based on the principles of natural selection, and are used to create virtual entities that are specialized for a particular task. Sets of these sorts of virtual machines might be more appropriately referred to as communities and populations, rather than as subroutines and programs.

The second of the three perspectives discussed here refers to the *computational abilities* of a biosystem. These arise from the operation of the virtual machinery described previously, and can be envisioned as forming an epistemic space of potential computational activities to which the mind is limited. Like virtual machines, computational abilities can be grouped by function into sets. Many researchers have proposed lists of candidate sets, or *faculties*, in order to delineate the mental architecture of naturally-occurring intelligent entities such as humans and other animals (Pinker, 1994; Gardner, 1993; Goldman, 1986). A similar taxonomy is proposed for the faculties of ecocyborgs, and is discussed later in this paper.

Finally, computation in biosystems can also be characterized by the information-processing activities that are actually performed. This movement through the space of potential computation is the dynamical manifestation of the computational abilities of a biosystem, and is referred to as *mentation*. (This is a general term that describes the activities of any biosystem; the term *thought* is used with reference to humans and similar animals.) Mentation can differ greatly between individual biosystems, in accordance with their goals, constraints, and unique experiences, even though their minds and computational abilities might be similar. For instance, two identical greenhouses might maintain entirely different internal climates in order to grow different species of plants.

The perspectives described here are useful when comparing the computational characteristics of biosystems that might be radically different in their physical structure and in their histories of experience and mentation. For instance, two ecocyborgs might differ enormously in their structure, the computational abilities of the first being based largely on virtual mechanisms

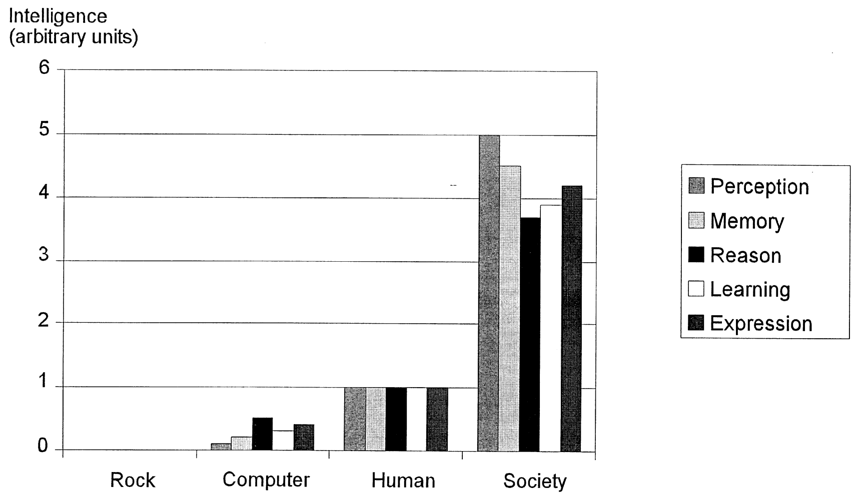

Fig. 2. Vectorized measures of intelligence, using an average human as the standard.

that are intrinsic to its ecosystem component, and extrinsic components forming the foundation for the mind of the second. The two entities could nevertheless have the same capacity to regulate their internal temperature in the face of climatic fluctuation. In the first ecocyborg the temperature regulation might be mediated by the thermal mass of a pond, whereas in the second this might be accomplished by a technological control network including thermostats, digital controllers, and propane heaters that are extrinsic to the ecosystem.

4. Intelligence

Intelligence measures are useful for comparing the computational abilities of different biosystems, and a variety of intelligence indices have been devised for use in various applications. In the past, for example, the mental ability of a human has often been viewed as a cohesive phenomenon, and has been characterized accordingly with a single-valued Intelligence Quotient. This is more informative than a binary distinction between *intelligent* and *not intelligent*, but an even more detailed description can be provided by evaluating a number of characteristics on continuous scales and then collecting their values into a vector. Strengths and weaknesses can then be compared among different biosystems if the scales are calibrated with standard points. Minsky (1985) suggested a scale of intelligence normalized in this way, for instance, with the mental ability of an average human defined as unity (Fig. 2). The adoption of such a scheme would be useful in the engineering of ecocyborgs with particular computational abilities, such as those required for autonomy. One basis for such a vectorized intelligence measure is the grouping of computational abilities into faculties. Accordingly, a set of faculties is proposed below for the particular case of ecocyborgs.

5. Mental faculties

In an extreme interpretation, the whole causal network that connects input with output can be considered as one, unified transfer function. Alternatively, an interpretation can be employed that distinguishes between types of computational abilities. Such a scheme inevitably results in indistinct categories that overlap to a degree, since in any taxonomy the manner in which computational abilities are grouped together is somewhat arbitrary. Some taxonomy must nevertheless be imposed in order to proceed with an analysis. Here, a scheme is presented that categorizes computational abilities into five groups: the faculties of perception, reason, memory, learning, and expression (Fig. 3). For each faculty there is a general discussion, which is illustrated with reference first to animals and then to ecocyborgs. The mental faculties of an ecocyborg might arise from either biological or technological components that could be either intrinsic or extrinsic to the ecosystem.

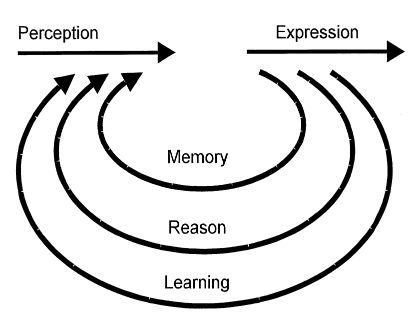

Fig. 3. The five mental faculties of an intelligent system.

5.1. Perception

The faculty of *perception* encompasses the ability to create information from signals. These signals may be of external origin, but if the biosystem is capable of self-observation then some may also originate internally. The abilities included in this faculty arise partly from virtual machinery embodied in an array of sensory devices (perceptors). The physical embodiment of the perceptors is of special relevance, since they form the interface with the physical surroundings. As well as creating information from incoming signals, the virtual machinery also transforms it so that it is accessible to other parts of the mind. If the information created by certain perceptors is always structured in a particular way, then the associated virtual machinery might be highly optimized for the specific tasks that are involved, as reflected by the intransigence of the physical substrate. Flexibility is sacrificed in this case, since the specialized configuration that results serves as a base-level filter for the information that is created.

In biological organisms such as mammals, the perception of external signals depends largely on massively parallel arrays of specialized sensory neurons in the epidermis, like the retinas of the eyes, the cochlea of the ears, and the olfactory buds in the nose and mouth. Specialized sensory neurons throughout the body also interpret internal signals. Highly adapted computational abilities are associated with each of these specialized arrays, which create information based on particular kinds of input signals. These kinds of abilities can dramatically impact the whole physiology and mode of existence of an entity. For example, bats have evolved to be extremely dependent on their ability for acoustic imaging, and the physiology of temperate plants is centered on the way that this type of vegetation perceives sunlight and seasonal changes in the environment.

In an ecocyborg, the ecosystem segment would have the inherent ability to perceive and respond to signals like solar radiation flux, rainfall, and the partial pressures of atmospheric gases. Technological mechanisms could also track these variables, as well as others that would not normally be perceived by a natural ecosystem, such as the unit cost of heating fuel. In the case of ecocyborgs with extensive ecosystem segments (intended for human habitation, for instance), it seems appropriate that any technological perceptor arrays should be massively parallel and highly distributed. This would result in the generation of large amounts of information, but because an ecocyborg would likely be immobile or primarily focused on managing its internal state, the task of perception would be somewhat simplified as compared to the case of an animal. Perception could be simplified even further if the internal sensors were immutable and immobile with respect to the rest of the biosystem.

5.2. Memory

Memory includes all of the abilities required to index, retain, and retrieve information. This can be interpreted as the ability to create or perceive patterns in information, or to create deeper semantic structures based on information generated through the faculty of perception. When new information is acquired, it is subsumed into the mind so that the structure of the constituent virtual machinery is contingent on the history of its mentation. This retained information is indexed by detecting any similarities to previous information. These relationships are made explicit through the creation of links between informational constructs, or equivalently, by grouping the constructs. This process is equivalent to the creation of a semantically deeper layer of information that can be described as *meta-information*. The indexing process can be iterated to create a richly structured network. The associative patterns within the network then serve to index the information and recall it in the appropriate context. Memory is therefore dependent on the capacity to detect, create, and compare patterns.

As with all mental phenomena, memory in animals arises from virtual machinery whose functioning corresponds to the physical activity of neurons. Animals with more developed nervous systems have correspondingly sophisticated memories that appear to correspond to the synchronized firing of many neurons (Greenfield, 1995). Patterns of relationships in retained information, i.e., associated memories, might correspond to the firing of subsets of neurons that are shared among various synchronized populations. Because of the vast numbers of neurons involved, it is possible to represent relatively large informational structures. The physiology of animal nervous systems has inspired the creation of similarly structured artificial neural networks. These have proven to be eminently capable of retaining, processing, and recalling patterns of information such as those that might be created by a biological sensory array.

In order to be significantly intelligent, ecocyborgs must retain, structure, and recall large amounts of information, just as animals do. The manner in which the required pattern-processing abilities will be implemented in ecocyborgs will depend on the underlying virtual machinery and the corresponding physical substrate in which it is embodied. Biological systems demonstrate an approach that involves massively parallel networks of information storage devices. In artificial systems, these devices might be packaged in a single structure, such as a silicon computer chip, but their basis will be ultimately reducible to large numbers of distinct components such as transistors. In order to support the required virtual machinery, these must be able to change state, and it should be possible to make their state dependent on that of other devices, so that they can be used to encode the sophisticated networks of information described previously. Finally, in order for this information to be kept current and accessible, there must be an interface with the other mental faculties of the ecocyborg.

5.3. Reason

It is speculated that increased autonomy improves the viability of an entity by heightening its ability to respond independently to an unpredictable environment. This implies the flexible and sophisticated formulation of appropriate responses to unforeseen stimuli. *Reason* is the faculty that encompasses the computational abilities required for this. It is bracketed by perception and expression, the faculties by which signals are translated into information and vice versa. Reasoning transforms the pool of information retained in the mind into mental products that potentially have an impact on the surrounding environment, or on the internal structure of the entity itself. These mental products include judgments, decisions, inferences, conclusions, and solutions to problems.

Human reason is the epitome of flexibility and sophistication, as evidenced by the great variety of behavior that it engenders. It is therefore difficult to completely catalog the abilities that it comprises, and it often seems that new ones become apparent in every scenario. A number of qualities can be used to characterize these abilities, corresponding qualities being definable to characterize the virtual machinery from which the abilities arise and the mentation that they enable. Of these, the qualities of mentation are the most commonly referred to. *Depth* and *breadth* are two of these: depth refers to the length of the chain of mediating events leading from inputs to outputs, whereas breadth indicates the number of alternative paths that are explored. Thus, reasoning might be fairly narrow and shallow, or it might simultaneously involve a vast array of different mechanisms in parallel and/or in series, each influencing the outcome to some degree. In the former case, the reasoning process might be sufficiently transparent so that the mentating system itself can observe, understand, and explain it; in the latter, it might be so diffuse and convoluted, with various branches reinforcing and inhibiting one another, that the process becomes entirely intractable. This is often referred to as *intuitive* reasoning. Reasoning may also be either *deductive* or *inductive*. Deduction moves from general premises to logical conclusions, and is supported by theoretical understanding, whereas induction is the inference of general principles from particular instances and relies on experience. Overall, the relationships between inputs and outputs can be extremely complicated, with many inputs taken into consideration and the activities of various reasoning mechanisms interacting with one another. The end result is often uncertain and multivalent.

If an ecocyborg is to have a high degree of autonomy, its mind must possess a wide variety of reasoning mechanisms that can interact flexibly with one another. The faculty of reason should therefore be composed of many semi-independent abilities that arise from such mechanisms, a scheme that is similar to some current interpretations of how the human mind functions (Pinker, 1997; Minsky, 1985). Each of these abilities could involve a different combination of the qualities mentioned above. The virtual machine that gives rise to each could operate on different kinds of information that might originate externally or be generated by other virtual mechanisms. The activity of this machinery might modify the internal state of the ecocyborg through the creation of mental products such as those mentioned above, and some of these mental products could also stimulate the faculty of expression to radiate signals into the surrounding environment.

5.4. Expression

The faculty of *expression* is the complement of perception. It encompasses the computational abilities required for the transposition of mental products into output signals. These signals can propagate outward to affect the external surroundings, or they can influence the internal state of the system. In a physical context, this involves the manipulation of material objects, whereas in a virtual setting it entails the manipulation of information, and can also include communication with other entities. As with perception, there can be one or more adjunct abilities permanently associated with each effector to enable the rapid and effective execution of habitual tasks, such as the parsing of directives intended for the effector.

As with the faculty of perception, some of the virtual machinery that underlies expressive ability forms an interface between the mind and the physical world, and so the physical embodiment of these virtual machines is again of particular relevance. In animals, effectors that impact the external surroundings are generally fewer in number and more localized than the vast arrays of perceptors described earlier. This is perhaps due to the tendency of a signal to disperse as it radiates from its source through an unconfined environment. The bulk of many animals is, nevertheless, made up of effectors and associated devices, through which physical signals are generated. For instance, the arms and legs of a human constitute effectors that interface with the external environment. There are also effector arrays, such as the peristaltic musculature, that influence the internal state of the body. Other expressive abilities, however, are oriented more toward the virtual rather than the physical realm, and so are not necessarily as directly dependent on the configuration of the material substrate in which they might be embodied.

As mentioned, most future ecocyborgs will probably be immobile, and therefore will not require the kinds of effectors that animals need for locomotion. External effectors will more likely be associated with activities such as maintaining a selectively permeable barrier between the ecocyborg and its surroundings, and with virtually-oriented tasks such as communication. Following the biological pattern, the internal effectors of an ecocyborg should be of a parallel and distributed nature, so that effects can be visited upon the entire extent of the system. Their type could vary greatly, depending on the nature of the ecocyborg; if it included a large ecosystem segment, the internal effectors could be as diverse as irrigation networks, air conditioning systems, or troops of pruning robots.

5.5. Learning

Learning includes the abilities that enable a mind to restructure itself adaptively. The idea of adaptation implies the improvement of performance, or increased viability in a particular context. Effective learning makes the mind of a biosystem more adept at interpreting the stimuli it encounters, and at responding in a manner that has favorable results. This requires that the biosystem be able to recognize in perceived information patterns that correspond to frequently encountered and exceptionally important environmental situations. The biosystem must also be capable of identifying associated patterns of mental activity that result in desirable outcomes in particular circumstances, and of generating new ones if the old ones are ineffective. In learning, important patterns are retained so that they can be quickly identified (in the case of perceived patterns) or reproduced (in the case of mental activities). The effectiveness of learning therefore depends on the ability to acquire or create new patterns and to retain those that are most useful. In a stable environment, this should make a biosystem increasingly successful, by whatever means this is measured. A changing environment could, however, require that the biosystem continuously restructure itself in order to deal with new situations. Depending on how challenging the environment is, a biosystem might not be able to keep pace, and it might become relatively less suited to its surroundings. There is more of an advantage if the faculty of learning is recursive, and can operate on itself to acquire better ways of learning. In a highly unpredictable environment, therefore, the autonomy of a biosystem is very dependent on its ability to learn, and on its ability to learn about learning.

Of all the biosystems that have been observed, humans are probably the most effective and versatile learners. Their ability to adapt to a wide variety of different environments is evidence of this. As suggested, the human faculty of learning encompasses the ability to adapt to significant environmental scenarios, and to determine which new scenarios are, in fact, significant. Humans can also reproduce courses of action that were successful in past circumstances, improve upon past actions, and, if necessary, even formulate entirely new strategies. Finally, humans can learn new ways of learning, indicating that this faculty can operate recursively on itself. For instance, a linguist who has learned several languages can draw upon past experience to acquire another one more quickly than someone who is unilingual.

In order to learn, an ecocyborg must be capable of recognizing, generating, evaluating, comparing, and reproducing patterns. The apparent ease with which the human mind accomplishes these tasks can be somewhat misleading. Cognitive scientists attempting to simulate these abilities on computers are discovering how difficult it is to reproduce them (Pinker, 1997). Nevertheless, methods have been developed that emulate some aspects of human learning, and that might also endow an ecocyborg with rudimentary learning abilities. One example is the training of artificial neural networks by backpropagation of error. An ecocyborg can only learn effectively, however, if it has the creative capacity to discover or invent new patterns of relationships. Creation in this context can involve *optimization*, whereby existing patterns are varied according to some scheme and the results are evaluated. More dramatic creative efforts are *exploratory*, involving variations that are radical departures from the established norm (Boden, 1990). Exploratory creation can proceed by *association*, where new relationships are established between two concepts in a kind of folding of *idea space*. In this way, previously disparate ideas are associated by identifying similarities between them, or transposing an idea from a familiar context to a new one. Finally, *inventive* creativity is the innovation of pattern in a foray into previously unexplored regions of idea space. Methods of implementing creative learning in ecocyborgs are speculative at this point, but a certain amount of consciousness would certainly increase the effectiveness of some associated activities, such as evaluating new phenomena or activities that directly involve the ecocyborg.

6. Consciousness

Although there is no universally accepted definition, consciousness is generally conceded to involve the ability to observe and reason about oneself. Based on this, a proposed working definition of *consciousness* is the maintenance by an entity of a self-referential model; i.e., a model that includes some representation of the entity itself, thus enabling it to reason about itself in relationship to its environment (Chalmers, 1990; Lacroix and Kok, 1991). The abilities that are necessary for consciousness in an ecocyborg are shared among all the mental faculties. Since consciousness is based on the creation of models, it requires, for instance, the perception of phenomena, the identification of patterns in the resultant information, and the creation of formal constructs that are similarly patterned.

The degree of consciousness of an entity can be measured on a continuous scale, as opposed to being regarded as a discrete, binary attribute. Human mentation, for instance, is sometimes deliberate, explicit, and transparent, but more often it is not directly observable by the reasoner himself. The human reasoner is therefore unable to generate a complete self-model, and is thereby less conscious than he might otherwise be. Although humans and many animals display various degrees of consciousness, natural ecosystems are only very slightly conscious by comparison, since they appear to lack the required abilities, virtual machinery, and corresponding physical substrates. It might be possible to make ecosystems more conscious however, by cyborging them with technological control networks.

Once a self-referential model has been generated it can be used in prediction, reflection, and imagination. *Prediction* is mentation about how real events might unfold in the future; *reflection* concerns how they developed in the past; and *imagination* deals with hypothetical alternatives to actual situations. Variations on this basic theme allow for more sophisticated mentation. The recursion of consciousness, for instance, involves the creation of models representing the entity in enough detail so that the existence of the self-referential model is also denoted. Accordingly, a model that provides an ecocyborg with a representation of itself, but from which any representation of consciousness is excluded, endows the ecocyborg with *primary consciousness* (Lacroix and Kok, 1991). An ecocyborg possesses *secondary consciousness* if the model does take itself into account, and so on for higher degrees of recursion. Ecocyborgs might also be engineered so as to be able to simultaneously instantiate a number of self-referential models, and so consciously reason in parallel about various problems and possible solutions. An ecocyborg that is able to reason consciously is likely to be more effective in its response to external phenomena than one that cannot do so. It would have a superior capacity to regulate its own internal state and to formulate appropriate external responses. This would increase its autonomy by making it more effective in the intentful pursuit of its own goals.

7. Autonomy

Autonomy is the independence of comportment that emerges when a sufficiently conscious mind can be described as possessing, to some degree, several defining characteristics (Kok et al., 1995; Bourgine and Varela, 1992). The first is *automation*: the capacity to operate without outside intervention. Although necessary, this alone is insufficient for significant autonomy, since even a clock, for example, is capable of indefinite operation without outside involvement. The second required characteristic is *volition*, or choice in action or thought. A highly automatic, volitive mind can respond to its environment in a flexible manner by defining its own goals and then formulating and executing strategies for attaining them. Advanced greenhouse control systems are being developed, for instance, that are capable of limited volition in fulfilling their operating requirements (Lacroix, 1994). Finally, in order to be significantly autonomous an entity must be *intentful*, and actually exercise its volition. Since the intentful pursuit of goals is involved, one could say that increased autonomy is equivalent to a greater degree of deliberate self-control (Conant and Ashby, 1970). In general, these goals minimally include the survival of the biosystem. In the case of engineered biosystems such as ecocyborgs, they could also include other design objectives.

Like intelligence and consciousness, autonomy should be measured on a continuous scale. Moreover, although autonomy is dependent on mind and consciousness, their presence to any extent is not in itself sufficient to ensure significant autonomy. Even a highly intelligent and conscious ecocyborg, for instance, could be extremely curtailed in its autonomy if it were engineered to pursue a very specific set of objectives, explicitly defining the necessary subgoals, and putting in place a rigid set of rules that governed its allowable attainment strategies. In contrast, an ecocyborg would be a great deal more autonomous if it were bound only by broad, long-term objectives and a loose set of guidelines. In calibrating such a continuous scale for autonomy, one might think it appropriate to use a theoretical maximum as a standard. This leads to a paradox, however, since complete independence in an entity requires a structure that is free of any implicit design objectives or behavioral biases that might influence the definition or pursuit of goals. The actual behavior of such an entity, moreover, would have to conform exactly to its intent, and not be influenced in any way by the environment. In the limit, therefore, absolute autonomy would require that the entity be responsible for creating itself as well as its external environment, and an absolutely autonomous system would therefore have to be absolutely creative. Since humans are incapable of imagining what such an entity might be like, it is difficult to use it as a calibration standard. The average of some human population could be used instead, as is often done for the calibration of scales of intelligence.

Although absolute creativity is an unattainable goal, any entity with some degree of autonomy must be creative enough to formulate at least a few of its own goals and behavioral guidelines. A significant degree of autonomy is desirable in any ecocyborg that is engineered to achieve particular goals in an unpredictable environment. An uncreative ecocyborg would be dependent on preformulated action plans that might not be suited to new situations, whereas a more creative one would be capable of adapting to unforeseen situations by restructuring its goal tree and implementing new strategies in order to achieve its overall objectives. An automated greenhouse, for example, could vary the parameter values of its regulatory models and simulations in order to optimize them for the current situation. More radical creative measures could be implemented in ecocyborgs that were faced with more challenging environments, but in order for them to be useful to their designers, their autonomy should be so shaped that they will not override their general design objectives.

8. Conclusions

Computers presently serve as the physical substrate for sophisticated virtual machinery that endows ecocyborgs with computational abilities that are superior to humans in some narrow domains. Such artificial constructs are, however, still vastly inferior to human minds in most computational tasks, and are completely incapable of performing others. As a result, the autonomy of existing ecocyborgs is very rudimentary, and they can operate without human supervision only under routine conditions. Some automated greenhouses can employ predictive control techniques to adapt to bounded fluctuations in feedstock quality or ambient temperature, for instance, but they cannot deal with large, unforeseen departures from normal operating parameters (e.g., Lacroix, 1994; Lacroix et al., 1996; Linker et al., 1998). In many circumstances it would be desirable to employ highly autonomous ecocyborgs that are capable of reasoning about themselves in the context of their environment, setting their own goals, devising strategies for their attainment, and executing them, all without human supervision. It is postulated that a high degree of autonomy is required of any unsupervised ecocyborg that must persist in an unpredictable environment.

The coherent lexicon and philosophy presented here facilitate the characterization and engineering of significantly autonomous systems, such as ecocyborgs. The creation of these will serve some practical purposes, but will also have an impact well beyond the utilitarian sphere. Highly autonomous ecocyborgs could be employed, for instance, to mediate the increasing human impact on extant natural ecosystems, and thus have a profound impact on human society. Entirely artificial ecocyborgs could also be created to serve a variety of other purposes, such as the production of food, fiber, and other biological products. Large, self-sufficient ecocyborgs could even provide a base for habitation and industrial expansion in space. Once proven technology has been developed for the construction of such entities, it may be possible to create them in great numbers, and perhaps even to make them capable of replicating themselves. In sufficient numbers, they might develop their own societies, collective structures that might evolve as computational systems in their own right, complete with economies, philosophies, and theologies. These societies might also be subject to engineering practices, in which case researchers can look ahead to shaping new structures not only at the level of individual ecocyborgs, but also at higher levels of conglomeration. The lexicon and philosophy provided here provide a language and framework with which such endeavors can be envisioned, planned, and executed.

Acknowledgements

The authors gratefully acknowledge the support of the Natural Sciences and Engineering Research Council of Canada, the J.W. McConnell Foundation, and McGill University. Appreciation is expressed to M.C. Stacey and P.C. Clark for reviewing drafts of this article.

References

Boden, M.A., 1990. The Creative Mind: Myths and Mechanisms. Basic Books, New York.

Bourgine, P., Varela, F.J., 1992. Towards a practice of autonomous systems. In: Varela, F.J., Bourgine, P. (Eds.), Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, 1st ed. MIT Press, Cambridge, MA, pp. xi–xvii.

Chalmers, D.J., 1990. Consciousness and cognition. On AI CD-ROM, Revision 2, 1993/1994. Network Cybernetics Corporation, Irving, TX. Available online at http://ling.ucsc.edu/~chalmers/papers/c-and-c.html.

Conant, R.C., Ashby, W.R., 1970. Every good regulator of a system must be a model of that system. International Journal of Systems Science 1, 89–97.

Dyson, G., 1997. Darwin Among the Machines. Addison-Wesley, Redwood City, CA.

Engelberg, J., Boyarsky, L.L., 1979. The noncybernetic nature of ecosystems. The American Naturalist 114, 317–324.

Gardner, H., 1993. Frames of Mind, 2nd ed. BasicBooks, HarperCollins Publishers, New York.

Goldman, A.I., 1986. Epistemology and Cognition. Harvard University Press, Cambridge, MA.

Greenfield, S.A., 1995. Journey to the Centers of the Mind. Freeman, New York.

Kok, R., Parrott, L., Clark, O.G., Molenaar, R., Lacroix, R., 1995. The engineering of autonomous biosystems: the EcoCyborg Project. CSAE Paper No. 95-602. CSAE, Saskatoon, SK, Canada.

Kok, R., Lacroix, R., 1993. An analytical framework for the design of autonomous, enclosed agroecosystems. Agricultural Systems 43, 235–260.

Lacroix, R., Kok, R., Clark, O.G., 1996. Use of a mulitasking operating systems as a setting for the simulation of an enclosed agroecosystem under cognitive control. Canadian Agricultural Engineering 38, 129–138.

Lacroix, R., 1994. A framework for the design of simulation-based greenhouse control. Ph.D. Dissertation, Department of Agricultural Engineering, McGill University, Montréal.

Lacroix, R., Kok, R., 1991. A cognitive controller with recourse to simulation. In: Automated Agriculture for the 21st Century ASAE Symposium, Chicago, IL, December.

Linker, R., Seginer, I., Gutman, P.O., 1998. Optimal CO_2 control in a greenhouse modeled with neural networks. Computers and Electronics in Agriculture 19, 289–310.

McNaughton, S.J., Coughenour, M.B., 1981. The cybernetic nature of ecosystems. The American Naturalist 117, 985–990.

Minsky, M.L., 1985. The Society of Mind. Simon and Schuster, New York.

Parrott, L., Clark, O.G., Kok, R., Molenaar, R., 1996. The EcoCyborg Project: new tools for the biosystems engineer. NABEC Paper No. 9655. Presented at the Northeast Agricultural and Biological Engineering Conference, Canaan Valley, WV, August 4–7.

Patten, B.C., Odum, E.P., 1981. The cybernetic nature of ecosystems. The American Naturalist 118, 86–895.

Pinker, S., 1994. Mind design. In: The Language Instinct, 1st ed. William Morrow, New York, pp. 404–430.

Pinker, S., 1997. How the Mind Works. W.W. Norton, New York.

Sagan, C., 1994. Pale Blue Dot: A Vision of the Human Future in Space. Random House, New York.

Vigilant, L., Stoneking, M., Harpending, H., Hawkes, K., Wilson, A.C., 1991. African populations and the evolution of human mitochondrial DNA. Science 253, 1503–1507.

Grant Clark received his undergraduate training in Agricultural Engineering at the University of Alberta (1993), and is currently completing his doctoral studies in the Department of Agricultural and Biosystems Engineering at McGill University.

Dr. Robert Kok, a Professor in the same department, received his B.E.Sc. in Chemical Engineering (1969) and his Ph.D. in Biochemical Engineering (1974) from the University of Western Ontario. His research interests include the application of computer-based technologies to environmental and agricultural issues, and the commercial uses of insects.

Dr. René Lacroix received his M.Sc. in Environmental Sciences from the Université du Québec à Montréal (1988) and his Ph.D. in Agricultural and Biosystems Engineering from McGill University (1994). He is presently employed as an Assistant Professor in the Department of Animal Science of McGill University. His current research involves the application of information technologies to dairy cattle breeding.