to appear in: Foundations of Science vol. 4 (2000)

Foundations and Methodology for an Evolutionary World View: a review of the Principia Cybernetica Project

Francis HEYLIGHEN

CLEA, Free University of Brussels, Pleinlaan 2, B-1050 Brussels, Belgium
E-mail: fheyligh@vub.ac.be

ABSTRACT. The Principia Cybernetica Project was created to develop an integrated philosophy or world view, based on the theories of evolution, self-organization, systems and cybernetics. Its conceptual network has been implemented as an extensive website. The present paper reviews the assumptions behind the project, focusing on its rationale, its philosophical presuppositions, and its concrete methodology for computer-supported collaborative development. Principia Cybernetica starts from a process ontology, where a sequence of elementary actions produces ever more complex forms of organization through the mechanism of variation and selection, and metasystem transition. Its epistemology is constructivist and evolutionary: models are constructed by subjects for their own purposes, but undergo selection by the environment. Its ethics takes fitness and the continuation of evolution as the basic value, and derives more concrete guidelines from this implicit purpose. Together, these postulates and their implications provide answers to a range of age-old philosophical questions.

Introduction

The Principia Cybernetica Project (PCP) is an international organization, founded by V. Turchin and C. Joslyn in 1989. They were joined in 1990 by F. Heylighen, to form the project's present "editorial board". The project's aim was to develop a complete scientific and philosophical world view, inspired by the theories of evolution, self-organization, systems and cybernetics. This collaborative development was to be supported by the most recent information technologies, such as the Internet, hypermedia, electronic discussion lists and semantic networks. Thus, the project proposed a unique integration of philosophical questions, scientific concepts and theories, and technological tools and methods. In spite of these very ambitious aims, a novel, untested methodology, and a mostly volunteer effort with very limited funding, PCP has managed to realize most of its objectives in the decade since it was founded.

Let us summarize its most important achievements. PCP created its own electronic discussion list (PRNCYB-L, recently renamed to PCP-discuss) in 1991, and its website, Principia Cybernetica Web (http://pcp.vub.ac.be/ or http://pcp.lanl.gov), in 1993. The website was intended to electronically publish the results of the project, in an interactive, semantically organized format. At present, Principia Cybernetica Web contains some 2000 cross-linked documents, which are consulted every day by thousands of people. Many of its features, such as a search function, menu bar, table of contents, automated "what's new" page, which were novel when they were introduced during the early days of the World-Wide Web, have now become commonplace. Others, such as a universal mechanism for creating public annotations to pages, are still rare, while its most ambitious innovation, a linking pattern that self-organizes, adapting to its users (Bollen & Heylighen, 1996), is as yet unique. Principia Cybernetica Web has become by far the most authoritative site in the broad domain of cybernetics, systems theory, complexity and evolution, as shown by its consistent high ranking in subject directories and search engines, the many references that are made to it, and the enthusiastic reactions of its users.

The content of Principia Cybernetica Web covers all the main components of a world view, as defined by Aerts et al. (1994). Its major subdivisions include: methodology, concepts, principles, history and future of evolution, metaphysics, epistemology, and ethics. As summarized in the following sections, this conceptual network proposes answers to most of the fundamental philosophical questions that people have been asking throughout the ages. The individual pages (called "nodes" of the conceptual network) discuss the most diverse topics: the evolution of cooperation, the mathematical concept of infinity, the cybernetic concept of control, selection criteria for knowledge, the origin of multicellular organisms, the second law of thermodynamics, the market mechanism, the principle of parsimony, etc.

This diverse material is structured quasi-hierarchically, which means that every "node" but the top node (home page) has one (or rarely more) "parent node", which is hierarchically superior to it. Moreover, all nodes are cross-linked in the fashion of hypertext, with concepts used in the discussion of a given topic pointing directly to the page where they themselves are discussed. The combination of this flexible hypertext structure with a well-thought out, logical arrangement of the different topics and arguments, has made it possible to organize this surprising variety of subjects in a more or less systematic and coherent way.

Of course, in spite of this extensive coverage, the Principia Cybernetica world view is far from complete, and every week new nodes or improvements to existing nodes are added. Obviously, a project as ambitious as the creation of an evolutionary, cybernetic world view will never be really finished. There will always remain issues that need further details and clarifications. Many of the pages are therefore in a kind of "half-ready" state, awaiting further streamlining and expansion. This makes the whole into a dynamic, constantly evolving system of concepts.

The ideas collected in Principia Cybernetica Web are obviously also not purely the product of the Project's efforts. Most of the tenets of the Principia Cybernetica philosophy already had been formulated before the start of the project, in the research of the projects' editors, V. Turchin (1977), C. Joslyn and F. Heylighen, in the work of other project contributors, such as D. T. Campbell (1974a, 1974b, 1979) and J. de Rosnay, or in the legacy of the visionary scientists, such as Ashby (1956), von Bertalanffy (1968), von Foerster (1979), and Wiener, whose innovations inspired the project. There is room for debate in how far the project's results are genuinely new, or merely the collection and organization of existing ideas. What is really novel, though, is the integration of this wealth and diversity of material into a practically and conceptually coherent system.

In spite of the noted limitations, the achievement is impressive, especially when compared to other attempts at interdisciplinary integration, which, in spite of the best intentions, often fail miserably. The present paper will explore some of the reasons for this success, focusing on the rationale behind the project, its scientific and philosophical assumptions, and the way these have influenced its practical implementation. Although the size and interconnectedness of the PCP conceptual network can only be appreciated by an extensive exploration of its website, this paper will thus attempt to review the most fundamental ideas behind the project.

Motivation of the Principia Cybernetica Project

The need for integration

It is a common observation that our present culture lacks integration: there is an enormous diversity of "systems of thought" (disciplines, theories, ideologies, religions, ...), but they are mostly incoherent, if not inconsistent, and when confronted with a situation where more than one system might apply, there is no guidance for choosing the most adequate one. Philosophy can be defined as the search for an integrating conceptual framework or "world view" (Aerts et al., 1994), that would tie together the scattered fragments of knowledge which determine our interaction with the world. Since the 19th century, philosophy has predominantly relied on science (rather than religion) as the main source of the knowledge that is to be unified.

After the failure of logical positivism and the mechanistic view of science, only one approach has made a serious claim that it would be able to bring back integration: the *General Systems Theory* (von Bertalanffy, 1968; Boulding, 1956). Systems theorists have argued that however complex or diverse the world that we experience, we will always find different types of *organization* in it, and such organization can be described by principles which are independent from the specific domain at which we are looking. Hence, if we would uncover those general laws, we would be able to analyse and solve problems in any domain, pertaining to any type of system.

Many of the concepts used by system theorists came from the closely related approach of *cybernetics*: information, control, feedback, communication... In fact, cybernetics and systems theory study essentially the same problem, that of organization independent of the substrate in which it is embodied. Insofar as it is meaningful to make a distinction between the two approaches, we might say that systems theory has focused more on the structure of systems and their models, whereas cybernetics has focused more on how systems function, that is to say how they control their actions, how they communicate with other systems or with their own components, ... Since structure and function of a system cannot be understood in separation, it is clear that cybernetics and systems theory should be viewed as two facets of a single approach. In order to simplify expressions, we will from now on use the term "cybernetics" to denote the broad domain of "cybernetics and general systems theory".

The fundamental concepts of cybernetics have proven to be very powerful in a variety of disciplines: computer science, management, biology, sociology, thermodynamics... A lot of recently fashionable approaches have their roots in ideas that were proposed by cyberneticians decades ago: complex adaptive systems, autonomous robots, artificial life, neural networks, human-machine interfaces, self-organization theories, systems therapy, organizational learning, ... Most of the fundamental concepts and questions of these approaches had already been formulated by cyberneticians such as Ashby, von Foerster, Bateson, McCulloch and Pask, ... in the 1940's and 1950's. Yet, cybernetics itself is no longer fashionable, and the people working in those new disciplines seem to have forgotten their cybernetic predecessors (Joslyn & Heylighen, 1999).

What is the reason that cybernetics does not seem to get the popularity it deserves? What distinguishes cyberneticians from researchers in the previously mentioned areas is that the former stubbornly stick to their objective of building general, domain-independent theories, whereas the latter focus on specific applications: data mining, psychotherapy, management, robotics... These applications attract attention insofar that they are useful, concrete or spectacular. On the other hand, the aim of general integration remains too abstract, and is not sufficiently successful to be really appreciated.

But why then is cybernetics less successful than these more trendy approaches? Clearly the problem of building a global theory is much more complex than any of the more down-to-earth goals of the fashionable approaches. But we may also say that the generality of the approach is dangerous in itself if it leads to remaining stuck in abstractions, which are so far removed from the everyday world that it is difficult to use them, interact with them, test them on concrete problems, in other words, get a feel on how they behave and what are their strengths and weaknesses.

The Principia Cybernetica Project started from the contention that the goal of global integration is still essential, but that cybernetics has a number of lessons to learn from its more specialized applications. Whereas cybernetics aims to unify science, it is in itself not unified. Instead of looking down on practical applications, it should try to understand how those applications can help cyberneticians in their task of unifying science, and first of all unifying cybernetics. It should look upon them as *tools*, that can be used for tasks that may extend much further than the ones they were originally designed for.

From Principia Mathematica to Principia Cybernetica

A situation similar to the one we sketched here arose around the end of the 19th century in mathematics. Mathematics proposed a great variety of very successful applications: geometry, calculus, algebra, number theory, etc. Yet, there was no overall theory of mathematics: these different domains functioned mainly in parallel, each with its own axioms, rules, notations, concepts, ... Most mathematicians would agree intuitively that these different subdisciplines had a "mathematical way of thinking" in common, but one had to wait for the development of mathematical logic by Boole and Frege, and set theory by Cantor before this way of thinking could be formulated more explicitly. Yet set theory and formal logic were still plagued by incoherence, paradoxes, inconsistencies and lacking connections.

One had to wait further for the classic work of Whitehead and Russell (1910-1913), the *Principia Mathematica,* in which they ground the "principles of mathematical thinking" in a clear, apparently consistent and complete way. (the theorem of Gödel later shattered the hope that such a formal treatment could ever be considered complete, but that is another story). What was novel in the work of Russell and Whitehead was that they applied mathematical methods to the foundation of mathematics itself, formulating the laws of thought governing mathematical reasoning by means of mathematical axioms, theorems and proofs. This proved highly successful, and the Principia Mathematica stills forms the basis of the "modern" mathematics as it is taught in schools and universities.

PCP is based on the assumption that something similar can be done with cybernetics: integrating and founding cybernetics with the help of cybernetical methods and tools. Similar to the mathematical application domains (number theory, geometry, etc.), the applications of cybernetics (neural networks, systems analysis, operations research, ...) need a general framework to integrate them. Similar to the integrating theories of mathematics at the beginning of the 20th century (Cantor's set theory, formal logic, ...), the integrating theories of cybernetics at the beginning of the 21th century (general systems theory, second-order cybernetics, ...) are not integrated themselves. In reference to Russell and Whitehead, our plan to develop unified foundations for cybernetics has therefore been called the "Principia Cybernetica Project" (Turchin, Joslyn and Heylighen, 1990; Joslyn, Heylighen & Turchin, 1993)).

Let us first explain what we mean by "unified foundations" for cybernetics, and thus for systems of thought in general. The PCP framework can be viewed as a philosophical system: that is to say a complete "world view" (Aerts et al., 1994), which is clearly thought out and well-formulated, avoiding needless ambiguity, inconsistency or confusion. It should integrate all the different domains of human knowledge, experience, and action. It should provide an answer to the basic questions: "Who am I? Where do I come from? Where am I going to?" Like in traditional philosophy it must contain at least an ontology or metaphysics (a theory of what exists in the world and where it comes from), an epistemology (a theory of how we can know the world around us), and an ethics or system of values (a system of goals and rules that can guide us in our actions).

Let us further indicate the similarities and differences between a *Principia Mathematica* and a *Principia Cybernetica*. Both mathematics and cybernetics are in the first place metadisciplines: they do not describe concrete objects or specific parts of the world; they describe abstract structures and processes that can be used to understand and model the world. In other words they consist of models about how to build and use models: metamodels (Van Gigh, 1986). This meta-theoretical point of view is emphasized in particular in the so-called “second order cybernetics” (von Foerster, 1979), which studies how observers construct models.

It is because of this metatheoretic character that mathematics and cybernetics can be applied to themselves: a metamodel is still a model, and hence it can be modelled by other metamodels, including itself (Heylighen, 1988). In mathematics, the best known illustration of such a self-representation is the technique of "Gödelization", where a proposition about natural numbers is represented by a natural number. Of course, it is well-known that any self-representation must be incomplete (Löfgren, 1990), as illustrated by the Gödel theorem, but we do not consider completeness in the formal sense to be a necessary condition for a practically functioning modelling framework.

Let us proceed with the differences between cybernetics and mathematics. Mathematics is characterized by the following assumptions: simplicity, regularity and invariance; the separability of systems into independent elements; and the objective, context-independent, value-free character of models. Cybernetics, on the other hand, emphasizes complexity, variety and process; the fact that elements only exists through their relations and interactions with other elements; and the subjective, context-dependent and value-dependent nature of models. Cybernetics does not deny the value of mathematics; it assumes it but goes beyond it, by trying to encompass phenomena which cannot be represented in a static, unambiguous, formal framework. It is clear then that the self-application of cybernetics, in the form of a Principia Cybernetica, must be different from the Principia Mathematica model. A Principia Cybernetica must put the emphasis on evolution and open-endedness, on different levels of precision or vagueness, on dynamic interactions between a variety of systems and viewpoints.

Part of the reason why the General Systems movement in the 1950's and 1960's did not succeed in its objectives was because its models and methods were still too dependent on the static, atomistic paradigm that gave birth to mathematics and classical mechanics. This led to a counterreaction in the 1970's and 1980's, second-order cybernetics, which emphasized the involvement of the subject or observer in any modelling process. Unfortunately, the discussions grown out of this *in se* important insight had a tendency to get stuck in endless recursions of the type of observers observing observers observing, and other forms of navel-gazing, thus losing touch with concrete applications. The reason why the present situation is much more promising is that we now not only have better concepts and insights, particularly about the evolutionary mechanisms that produce complex systems, but also better practical tools (e.g. computers, the Internet), and methods for modelling complex and dynamic phenomena.

Although we must take into account the intrinsic complexity and context-dependence of the phenomena we wish to model, the idea of developing general principles (Principia) still assumes a striving towards clarity, "objectivity", and invariance. We should avoid getting trapped in endless discussions and confusions over subjective meanings and viewpoints. The invariant principles that can be derived (see e.g. Heylighen, 1992b), however, are situated at such a high level of abstraction that they do not impose absolute restrictions on concrete issues. They rather form an empty skeleton or framework, on which a multiplicity of more concrete theories can be hung (cf. Boulding, 1956). This framework can function primarily as a heuristic tool for building models and suggesting hypotheses. It will not *a priori* exclude any theory, but provide guidelines for avoiding ill-formulated problems and conceptual confusions, and for improving existing models. To achieve this, the framework must incorporate methods for concretizing its own recommendations, depending on the context of the problem. This means that, unlike mathematics, the framework should provide many intermediate levels between the abstract, precise and invariant principles and their concrete, context-dependent implementations.

Let us now summarize the basic concepts and principles that form the foundation for the Principia Cybernetica world view. This cybernetic philosophy, which forms the core of the project's theoretical content, can then be used as inspiration to suggest methods and tools for the project's practical implementation.

An evolutionary-cybernetic philosophy

Systems and evolution

The philosophy that provides the starting point for the Principia Cybernetica is basically *systemic*, like in the original ideas of the founders of General Systems Theory (von Bertalanffy, 1968; Boulding, 1956). This means that we assume that phenomena cannot be understood in isolation: they must always be viewed in relation to other phenomena. Networks of such relations may define "systems": basic forms of organization with a relatively stable identity. This identity is typically supported by some form of "organizational closure": an organization that produces and maintains itself, thus defining itself as an autonomous entity (Joslyn, 2000a; Varela, 1979). This can be illustrated by the negative feedback loop that underlies a cybernetic control system, such as the thermostat, or by an autocatalytic cycle of chemical reactions that is supposed to be at the basis of the origin of life (Kauffman, 1993).

Such organization is independent of the material substrate or components of the system: the same organization can in principle be implemented as a network of neurons, transistors, molecules or people. Hence it can be modelled by abstract structures and properties as proposed in mathematics and information theory. It is this inherent substance- or domain-independence that allows us to represent disparate systems in the same way, whether they be physical, chemical, biological, psychological, social, technological or mathematical.

Organization is not just structure, it also entails function: systems are in general characterized by goals that they try to reach through specific strategies and schemes. This goal-directedness is implemented through mechanisms of control (Powers, 1973; Ashby, 1956; Turchin, 1977), such as feedback and feedforward, which rely on communication between different systems or subsystems. Systems are basically dynamic, adapting their behavior to different circumstances and objectives.

In addition to these traditional systemic assumptions PCP starts from the principle of *evolution*: systems are not given, fixed organizations, they come into being, develop, disappear, or change into different systems. Systems are the result of a continuing process during which more and more complex forms of organization emerge. This evolution itself does not have a final goal, but is directed by the trial and error process of natural selection. Different (re)combinations of systems are formed by variation, but only those combinations are retained that are "fit", i.e. stable and/or capable of fast (re)production (Heylighen, 1999. The stability of the organization is what turns the combination or assembly into a "system". The variation process may be guided by knowledge acquired earlier, but in its most basic form it is *blind*: it does not know where it is going, or which of the variants it generates will be selected (Campbell, 1974a).

This evolutionary interpretation may also be called *constructive*: it assumes that systems can only be really understood by analysing the process through which they have been assembled. The variation and selection mechanism continuously constructs new systems from previous, usually simpler, systems. These building blocks themselves have emerged from even simpler components, which are the result of combinations of yet more primitives parts, ... The properties of the system cannot be reduced to the properties of their components: they can only be understood as results of the construction process itself, of the specific way in which the components have been assembled in a stable, selectively retained organization.

This constructive philosophy is related to constructive mathematics, where the existence of an object can only be proven by explicitly retracing its construction process, and not by merely disproving its non-existence on the assumption of the principle of the excluded middle. It is also related to constructivism in epistemology: the thesis that knowledge is not a passive apprehension or reflection of external reality, but an active construction by the subject.

Constructivism and foundationalism

In the limit, such a constructive philosophy entails that one cannot be satisfied by a philosophy that is based on "fundamental laws of nature", "first causes" or "prime movers", that is to say on fixed foundations which cannot be analyzed further. Whatever principle or organization is at the base of a construction process, it is itself merely the result of a previous construction and hence cannot in any way be ultimate. The only "primitives" that can be accepted in a constructive philosophy must be so simple as to be *empty of organization* of any form. All others, including the fundamental laws of physics, are to be viewed as the result of evolution through variation and selection, and must be analysed as such. An example of such an "empty" fundamental is the *tautological principle of natural selection*: fit systems survive, unfit systems are eliminated (Heylighen, 1992b, 1999).

A constructive philosophy is *anti-foundational*, in the sense that it rejects any contentful, fixed foundations. For example, traditional empiricist epistemology considers perception as a foundation for the theory of knowledge, with perceptual data as the building blocks out of which knowledge systems must be built. However, evolutionary epistemology rejects such a perceptual foundationalism and analyses perception as just one of the many selectors determining the adaptation process of the system to its environment (Campbell , 1974).

Yet a constructive philosophy can be considered foundational in the sense that it takes the principle of constructive evolution itself as a foundation. This principle is different from other foundations, however, because it is empty (anything can be constructed, natural selection is a tautology), but also because it is situated at a higher, "meta" level of description. Indeed, constructivism allows us to interrelate and intertransform different foundational organizations or systems, by showing how the one can be constructed from the other, or how two different systems can be reconstructed from the same, more primitive organization. In order to clarify this further we will have to introduce the concepts of emergence and metasystem transition.

A process metaphysics

Philosophies traditionally start with an ontology or metaphysics: a theory of being in itself, of the essence of things, of the fundamental principles. Examples of ontologies are proposed by Newtonian mechanics, which sees hard, elementary particles moving in space according to deterministic "laws of nature" as the essence of the world, and by the traditional monotheistic religions, which see the world as created and governed by the Will of God. In a traditional systemic philosophy "organization" might be seen as the fundamental principle of being, rather than God or elementary particles. In a constructive systemic philosophy, on the other hand, the essence is the *process* through which this organization is created.

There have been several attempts at building a process metaphysics, by philosophers such as Bergson, Whitehead (1929) and Teilhard de Chardin (1959). However, these early process philosophies are characterized by vagueness and mysticism, and they tend to see evolution as goal-directed, teleological, guided by some supraphysical force, rather than as the blind variation and selection process that we postulate. They are thus not constructivist in the radical sense as defined above.

The ontology we propose starts from elementary actions or processes (Turchin, 1993b), rather than from static objects or particles. An elementary action is for example the decay of a radio-active particle into a number of smaller particles, or the creation of a molecule by the bonding of two smaller molecules or atoms. A process is then merely a sequence of such actions. Complex systems are constructed by such processes through the mechanism of random combination of components, and the selective retention of stable combinations (Simon, 1962). This leads to a *self-organizing* evolution of the universe as a whole.

It is characterized by the spontaneous emergence of ever more complex organizations (cf. Heylighen, 1999; Heylighen, Joslyn & Turchin, 1995; Turchin, 1977) during evolution: from space-time and elementary particles, to atoms, molecules, crystals, dissipative structures, cells, plants, animals, humans, society, culture... In this hierarchy of system types (Boulding, 1954), the phenomena studied by cybernetics typically start from about the level of thermostats or dissipative structures. Yet, a process ontology can also be used at a much lower level, for example to reconstruct the elementary structures of space-time (Heylighen, 1990a), or the fundamentals of set theory (Turchin, 1987). A reconstruction of the most important stages in this global evolution allows us to answer the questions: "Where do I come from? Who am I?"

Processes of emergence are the "quanta" of evolution: discontinuous transitions which do not change just the state of a system but its organization itself. They lead to the creation of a new system with a new identity, obeying different laws and possessing different properties (Heylighen, 1991a). In such a system, the behaviour of the whole is constrained by the parts (a "reductionist" view), but the behaviour of the parts is at the same time constrained by the whole ("downward causation", i.e. a "holistic" view) (Campbell, 1974b).

Perhaps the most important type of emergence is the "*metasystem transition*" (Turchin, 1977; Heylighen, Joslyn & Turchin, 1995; Heylighen, 1995), which results in a higher ("meta") level of *control* while increasing the overall *freedom* and *adaptivity* of the system. Examples of metasystem transitions are the emergence of life, multicellular organisms, the capacity of organisms to learn, and human intelligence (see also Maynard-Smith & Szathmary, 1995). A metasystem transition is characterized by an increase of variety at the object level (usually through the assembly of a multiplicity of object systems), together with the emergence of a control at the metalevel, which coordinates, and chooses from, the variety of actions available at the level below (Heylighen, 1991a, 1995). In this view, evolution is fundamentally progressive (Heylighen, 1999; Heylighen & Bernheim, 2000b Stewart, 2000): it moves preferentially in the direction of increasing fitness, and metasystem transitions can be seen as discrete steps in this on-going progress towards more synergetic and adaptive systems (Turchin, 1977).

The transition from a foundational to a constructive philosophical framework can be seen as an example of a metasystem transition: instead of starting from a single set of foundations, all possible foundations are taken together (assembling a multiplicity of object systems), and a meta-foundational control (the principle of constructive evolution) is asserted as a way to integrate them.

A constructive epistemology

This example shows that a metasystem transition is not only a mechanism of physical or biological evolution, but also a basic problem solving principle. In fact evolution itself can be likened to a gigantic problem-solving process searching through trial and error for an answer to the question: how to build a system that will survive in a maximal variety of situations? (Heylighen, 1999) Knowledge or cognition is one of the results of that search: a mechanism that makes systems more efficient in surviving different circumstances, by short-cutting the purely blind variation and selection (Campbell, 1974). Knowledge functions as a *vicarious selector* which selects possible actions of the system in function of the system's goal (ultimately survival) and the situation of the environment. By eliminating dangerous or inadequate actions before they are executed, the vicarious selector foregoes the selection by the environment, and thus increases the chances for survival of the system.

A vicarious selector can be seen as the most basic form of a *model*: an abstract system representing processes in the environment (Turchin, 1993a; Joslyn, 2000b). A model is necessarily simpler than the environment it represents, and this enables it to run faster than, i.e. *anticipate*, the processes in the environment. It is this anticipation of interactions between the system and its environment, with their possibly negative effects, that allows the system to compensate perturbations before they have had the opportunity to damage the system. Models function as recursive generators of predictions about the world and the self: it is not necessary that a piece of knowledge produce an actual prediction, it is sufficient that it produces another piece of knowledge, which itself produces a piece of knowledge, ..., and so on, until the last piece effectively produces a prediction (Turchin, 1991, 1993a).

Models are not static reflections or homomorphic images of the environment, but dynamic constructions achieved through trial-and-error by the individual, the species or the society. This construction of models is similar to the continuous construction of systems by variation and selection that takes places everywhere in the universe. What models represent is not the *structure* of the environment but its *action,* insofar as it affects the system. They are both *subjective*, in the sense of being constructed by the subject for its own purposes, and *objective,* in the sense of being naturally selected by the environment: models which do not recursively generate adequate predictions are likely to be later eliminated. This can be understood from the interaction between a list of objective, subjective and intersubjective selection criteria for knowledge (Heylighen, 1997). There is no "absolutely true" model of reality: there are many different models which each may be adequate in solving particular problems, but no model is capable to solve all problems. The most efficient way to choose or to construct a model which is adequate for the given problem is by reasoning on a metacognitive level, where a class of possible models can be analysed and compared (cf. Heylighen, 1988). This requires a metasystem transition with respect to the variety of individual models.

An evolutionary ethics

This evolutionary philosophy can also be used to develop an ethics or system of values. The basic purpose here is the continuation of the process of evolution, avoiding evolutionary "dead ends" and minimizing the probability of extinction. Natural selection entails survival and reproduction, summarized in the concept of fitness (Heylighen, 1999), as the essential value.

The idea of an evolutionary ethics has not been very popular until now, and we will therefore go into a little more detail about this aspect of our philosophical system. Evolutionary ethics got a bad reputation because of its association with the "naturalistic fallacy": the mistaken belief that human goals and values are determined by, or can be deduced from, natural evolution (Campbell, 1979). Values cannot be derived from facts about nature: ultimately we are free in choosing our own goals (Turchin, 1991).

However, we must take into account the principle of natural selection, which implies that if our goals are incompatible with the conditions necessary for survival, then we will be eliminated from the natural scene. Of course, there is no natural law or absolute moral principle which forbids you to commit suicide, but you must be aware that this means that the world will continue without you, and that it will quickly forget that you ever have been there. If we wish to evade this alternative, this means that we will have to do everything for maximising survival.

A second fallacy to avoid is the naive extrapolation of past evolution into the present or future. The mechanisms of survival and adaptation that were developed during evolution contain a lot of wisdom—about past situations (Campbell, 1979).

They are not necessarily adequate for present circumstances. This must be emphasized especially in view of the creativity of evolution: the emergence of new levels of complexity, governed by novel laws.

For example, as shown by the theory of r/K selection in biological evolution, there is a trade-off between fast reproduction (r selection) and long life (K selection): organisms that reproduce quickly (e.g. mice, insects) generally do not live long and therefore do not have the time or the resources to become very large, strong, efficient or intelligent, as K-selected organisms would (e.g. elephants, tortoises, people); r-selection is selection for quantity rather than quality. Quick reproduction (r) is advantageous in environments where the probability of death is high, but where there are enough resources for a fast growing population. Continuing development (K) is preferable in an environment where the risks of unexpected death are under control, but that is limited in its carrying capacity in terms of available resources. Our present society becomes more and more like a K-environment, and therefore evolutionary theory would admonish us to increasingly focus on the values associated with K selection, i.e. long life and extended development, while suppressing the inherited tendency to produce a lot of offspring.

As a more radical example of a needed change in evolutionary strategies, biological evolution, based on the survival of the genes, has favoured selfishness and nepotism: maximizing one's own profit, with a disregard for others (unless those others carry one's own genes: close family). In a human society, on the other hand, we need moral principles that promote cooperation, curbing too strong selfishness. Once the social interactions have sufficiently developed the appearance of such moral principles (e.g. "thou shalt not steal") becomes advantageous, and hence will be reinforced by natural selection, even though it runs counter to previous "selfish" selection mechanisms (Campbell, 1979; Heylighen & Campbell, 1995). The development of human society is an example of a metasystem transition, which creates a new system evolving through a mechanism which is no longer genetic but cultural (Turchin, 1977).

One of the implications of that transition concerns the interpretation of survival. In biological evolution survival means essentially survival of the genes, not so much survival of the individuals (Dawkins, 1976). With the exception of species extinction, we may say that genes are effectively immortal: it does not matter that an individual dies, as long as his or her genes persist in off-spring. The death of individual organisms can even contribute to genetic fitness, by focusing resources on reproduction rather than individual survival, as illustrated by salmons dying shortly after they have spawned.

Individual death does not benefit cultural evolution, though. In socio-cultural evolution, the role of genes is played by cognitive systems ("memes", Dawkins, 1976; Heylighen, 1997), embodied in individual brains or social organizations, or stored in books, computers and other knowledge media. However, most of the knowledge acquired by an individual still disappears at biological death. Only a tiny part of that knowledge is stored outside the brain or transmitted to other individuals. Further evolution would be much more efficient if all knowledge acquired through experience could be maintained, in order to make place only for more adequate knowledge.

This requires an effective immortality of the cognitive systems defining individual and collective minds: what would survive is not the material substrate (body or brain), but its cybernetic organization. This may be called "cybernetic immortality" (Turchin, 1991). We could conceive its realization by means of very advanced man-machine systems, where the border between the organic (brain) and the artificially organic or electronic media (computer) becomes irrelevant. The death of a biological component of the system would no longer imply the death of the whole system.

Cybernetic immortality can be conceived as an ultimate goal or value, capable to motivate long-term human action. It is in this respect similar to metaphysical immortality (Turchin, 1991): the survival of the "soul" in heaven promised by the traditional religions in order to motivate individuals to obey their ethical teachings (Campbell, 1979), and to creative immortality (Turchin, 1991): the driving force behind artists, authors or scientists, who hope to survive in the works they leave to posterity.

Another goal that can be derived from the basic values of survival and development is "self-actualization" (Maslow, 1970; Heylighen, 1992a): the desire to realize all our human potential, that is to say, to maximally develop the knowledge, intelligence and wisdom which may help us to secure survival for all future contingencies. Self-actualization may be defined as an optimal, conscious use of the variety of actions we are capable to execute. Self-actualization is strongly correlated with happiness, subjective well-being, or general satisfaction with life. By empirically determining which social and economic conditions contribute most to happiness for the world population, it is possible to derive a more concrete list of universal values (Heylighen & Bernheim, 2000a). These include health, wealth, knowledge, safety, equality and freedom.

The main remaining problem of an evolutionary ethics is how to reconcile the goals of survival and development on the different systems levels: the level of the individual (personal freedom), of society (cooperation between individuals), and of the planet (survival of the world ecology as a whole). The necessary competition between levels follows from the problem of suboptimization, according to which what is best for a subsystem is not always best for the encompassing system (Heylighen & Campbell, 1995). In its stronger version, this problem takes the form of the "tragedy of the commons": the destruction of shared resources because of individually selfish optimization. It is clear that the different levels have very complicated interactions in their effect on selection (Campbell, 1979), and therefore we need a careful cybernetic analysis of their mutual relations.

Eternal questions: a summary of the PCP world view

The Principia Cybernetica world view is intended to answer the fundamental questions that every person reflecting about the world and his or her place in it has been asking throughout the ages. The PCP philosophy is organized as a complex network of mutually dependent concepts and principles. Therefore, the answers to these questions are scattered throughout the different "nodes" of the PCP website. The present section brings these different questions and answers together. The answers given here are by necessity short. They barely scratch the surface of a profound and complex issue. For more detail, the reader is referred to the Principia Cybernetica website, where the different answers are linked to more detailed expositions of their basic concepts.

What is? This question defines the domain of ontology. We believe that the fundamental stuff of being, the essence of the universe, consists of elementary processes or actions (Turchin, 1993b; Heylighen, 1990), rather than matter, energy or ideas. Complex organizations, such as atoms, molecules, space and time, living beings, minds and societies emerge out of these actions through the process of evolution.

Why is there something rather than nothing? The universe arose spontaneously, through self-organizing evolution, based on the self-evident principles of variation and natural selection (Heylighen, 1992b). Any possible variation (for example a "quantum fluctuation of the vacuum") would be sufficient to set the self-organizing process in motion, thus generating a complex universe with its diverse components and structures.

Why is the world the way it is? The specific state of the universe or the world in which we live is partially a historical accident, since evolution is an indeterministic process, partially the result of an orderly process of self-organization, which leads predictably to higher levels of organization through the mechanism of metasystem transition (Heylighen, 1999).

Where does it all come from? We can reconstruct in some detail the subsequent stages in the evolution of the universe, leading from the Big Bang, elementary particles, atoms and molecules to living cells, multicellular organisms, animals, people and society. Thus, the history of evolution, conceived as a sequence of metasystem transitions, tells us how and in which order all the different types of phenomena we see around us have arisen (Heylighen, 1995; Turchin, 1977).

Where do we come from? Humans evolved out of animals that had the capacity to learn associations from the environment, by additionally acquiring the capacity to think, that is, autonomously control these associations. Human thought is rooted in the emergence of symbolic language (Turchin, 1977; Heylighen, 1995).

Who are we? As far as we know, humans occupy the provisionally most advanced level in the hierarchy of metasystems. Our capacity for thought distinguishes us from the animals by giving us uniquely human characteristics, such as self-consciousness, imagination, planning, play, sense of humor and esthetic feelings (Turchin, 1977).

Where are we going to? The theory of metasystem transitions helps us to extrapolate present, on-going progress into the future. Recent developments point to a new metasystem transition which will bring us to a yet higher level of complexity or consciousness, transcending individual thought. This emergent level is perhaps best described by the metaphor of the social superorganism and its global brain (Heylighen & Bollen, 1996; Turchin. 1977).

What is the purpose of it all? Evolution does not have a purpose, in the sense of a fixed goal to which it is advancing (Teilhard's (1959) "Omega point"). However, although evolution is largely unpredictable, it is not random either. Selection can be seen as having the implicit goal of maximizing survivability or fitness. This implies a preferred direction of evolution (Heylighen, 1999), which is in practice characterized by increasing complexity, adaptivity and intelligence.

Is there a God? Since the principles of self-organizing evolution satisfactorily explain the origin and development of the universe, and our place in it, there is no need to postulate a personal God, in the sense of a conscious entity outside of the universe which created that universe. But since in our ontology every action can be assigned to, or labelled by, an agent (Turchin, 1993b), for those who wish to retain the concept of "God", it is possible to define this concept as the "agent" of the Big Bang. Similarly, it is possible to consider the universe as a whole, or the process of evolution itself, as God-like, in the spirit of pantheism.

What is good and what is evil? The evolutionary mechanism of natural selection makes an implicit distinction between "good" or "fit" situations (those which survive in the long term), and "bad" or "unfit" ones (those which are eliminated sooner or later). Therefore, we might equate good or higher values with anything that contributes to survival and the continuation of the process of evolution, and evil with anything that destroys, kills or thwarts the development of fit systems.

What is knowledge? This question defines the domain of epistemology. Knowledge is the existence in a cybernetic system of a model, which allows that system to make predictions, that is, to anticipate processes in its environment. Thus, the system gets control over its environment. Such a model is a personal construction, not an objective reflection of outside reality (Heylighen, 1997; Turchin, 1993a).

What is truth? There are no absolute truths. The truth of a theory is merely its power to produce predictions that are confirmed by observations (Turchin, 1993a). However, different theories can produce similar predictions without one of them being right and the other wrong. "True" knowledge is the one that best survives the natural selection for predictive power (Heylighen, 1997; Campbell, 1974).

What is consciousness? Even the most primitive cybernetic agent must be able to sense its environment in order to reach its goals. More complex agents can integrate such sensations into a global awareness of their situation and learn from these experiences. Only people are moreover conscious of their own experiences, and thus able to control them. Such consciousness is an dynamic, adaptive relation between the subject and its situation. It is not some mysterious "fluid" or "field", and there is no intrinsically "hard" problem of consciousness.

Do we have a free will? The indeterminacy of quantum mechanics has done away with the Newtonian world view in which all future events are predetermined. In our evolutionary world view, there always is a choice: a variety of possibilities, some of which are retained by selection. A cybernetic agent by definition has some degree of control over that selection. Because of their capacity for thought, people are not only free to choose between given possibilities, but able to conceive novel possibilities, and explore their consequences.

How should we act? Effective action is based on a clear sense of goals or values, and a good model of the environment in which you try to reach these goals. By applying problem-solving methods, you can explore the possible situations in your model to find the most efficient path from your present situation to your goal (Heylighen, 1988). You can then try out this action plan in practice, taking into account the feedback you get in order to correct your course.

How can we be happy? People are happy when they are "in control", that is, feel competent to satisfy their needs and reach their goals (Heylighen & Bernheim, 2000a). Happiness is most common in societies which provide sufficient wealth, health care, education, personal freedom and equality. Happy people tend to be self-confident, open to experience and have good personal relations. Promoting these social and personal values should increase our overall quality of life.

Why cannot we live forever? Evolution has predisposed us to age and die because fitness is achieved more easily by quick reproduction than by long life. Aging is the result of a variety of deterioration processes. Therefore, it is unlikely that we will achieve biological immortality in the near future, in spite of a constantly increasing life span. However, we can still aim for cybernetic immortality: survival of our mental organization, rather than our material body (Turchin, 1991).

What is the meaning of life? This question in a sense summarizes all previous questions. It is usually understood as: "What are the highest values, the Supreme Goals which I should try to achieve?". We stress that every human being must freely set those goals for himself or herself. The supreme goal which we choose derives logically from our cybernetic world view: to make a constructive contribution to the evolution of humanity, in order to maximize our long-term chances of survival (immortality). In essence, the meaning of life is to increase evolutionary fitness.

Methods and tools for developing a cybernetic philosophy

The methodology used by the Principia Cybernetica Project to build a complete philosophical system is based on a "bootstrapping" principle: the form through which the knowledge is expressed affects its content, and vice versa. Thus, our theories about the evolutionary development of systems are applied to the development of the theory itself, while the structuring of concepts in the form of an evolving semantic network suggests new theoretical concepts.

When constructing a philosophical system the fundamental building blocks that we need are *ideas*: concepts and systems of concepts. Ideas, similarly to genes, undergo a variation-and-selection type of evolution, characterized by mutations and recombinations of ideas, and by their spreading and selective reproduction or retention (Heylighen, 1997; Dawkins, 1976). The basic methodology for quickly developing a system as complex as a cybernetic philosophy consists in supporting, directing and amplifying this natural development with the help of cybernetic technologies and methods.

It requires, first, a large variety of concepts or ideas, provided by a variety of sources: different contributors to the project with different scientific and cultural backgrounds. Second, we need a practical tool for representing and manipulating these concepts: the computer. Third, we need a system that allows the representation of different types of combinations or associations of concepts: a hypermedia network. Fourth, we need selection criteria, for picking out new combinations of concepts, that are partly internal to the system, partly defined by the needs of people that are using the system. Finally, we need procedures for reformulating the system of concepts, building further on the newly selected recombinations, with the help of the concepts of emergence, and especially of metasystem transition. We will now discuss some practical tools for supporting each of these stages in the process (cf. Joslyn, 1991).

Supporting a variety of collaborators

The first difficulty with the practical development of a Principia Cybernetica is that something as large and as complex as a full cybernetic philosophy, integrating the fundamental ideas from all the different scientific disciplines, cannot be elaborated by a few individuals. The projects need a large variety of contributors from many different backgrounds (the traditional scientific disciplines, but also e.g. philosophy, technology, religion, art, ...), but with a common interests in the objectives of the project, to provide ideas, suggestions and feedback. These people are typically scattered over different countries and regions. The only medium that allows them to collaborate efficiently on this common project is the Internet.

The Principia Cybernetica Web is intended as a shared environment that allows contributors from all around the globe to add to a growing body of knowledge. Two types of collaborators are distinguished: editors and contributors (Heylighen, Joslyn, & Turchin, 1991; Joslyn, Heylighen, & Turchin, 1993). The editors have the final responsibility for the organization and selection of the material, like the editors of a book or a journal. Contributors can either be invited to write a specific section by the editors, or submit material on their own initiative. Although the editors have the right to reject material they deem inappropriate or of insufficient quality, the low cost of electronic publication allows us to be quite liberal. As long as there is sufficient storage space, all contributions can be published.

Another difference between PCP and a traditional scientific journal is that people can contribute even without submitting original material. This can happen by providing existing documents for electronic publication on Principia Cybernetica Web, such as the Glossary from the American Society for Cybernetics, or books that are out of print, such as Ashby's (1956) "Introduction to Cybernetics". Contributors typically also make suggestions about relevant work or interesting ideas produced by others, stimulating the editors to expand the conceptual network, or to point out connections/differences between the Principia Cybernetica philosophy and related approaches. Finally, contributors give constant feedback about the material that is already there, allowing the editors to correct mistakes, clarify points that tend to be misunderstood, or give more prominence to topics that many people are interested in.

Still, it is necessary to distinguish contributions fulfilling all the editorial standards from those that are more questionable, so as not to degrade the overall quality of the Web. The quasi-hierarchical organization of PCP Web allows an almost continuous gradation: the highest levels are the ones most easily reachable from the entry page, and are therefore most likely to be read by a large public. They should therefore be under strict quality control. For levels situated lower, which may only be reachable after an extended path, the criteria can be less strict: these nodes may contain material that is incomplete, not very clear or controversial. Since the Web is in a continual flux, one should be able to live with unfinished material, following the principle that it is better to have a mere first thought about an important problem, than to have nothing at all. A coarse and incomplete sketch may eventually become the centrepiece of an extensive conceptual framework. During this development, more and more links will be attached to it, so that it becomes more easily reachable and attains a more prominent place in the web. This upgrading of the position of a node in a hypertext can be done automatically, in a self-organizing manner, as will be discussed in the next section.

It is still important to make a clear separation between editorially selected nodes and others. We distinguish in principle three types: *consensus nodes*, on which all editors agree, *individual contributions*, which are considered to be within the general spirit of PCP but without each editor needing to agree about every phrase in the text, and *discussion nodes*, which are not under editorial control (Heylighen, Joslyn & Turchin, 1991; Joslyn, Heylighen & Turchin, 1993). The last category is distinguished from the former two by the way it is implemented: as an *annotation*. Everybody reading Principia Cybernetica Web can add his or her own comments, with criticisms, questions or elaborations, to any of the nodes. It suffices to click on the "Add Comment..." link. This calls up a form in which the user can enter his or her comment. These comments are immediately published to the web in the form of a new node, which is linked to the node commented upon by an entry in its “Discussion” field. At this moment, Principia Cybernetica Web has about 500 discussion nodes, and about the same number of editorially selected ones, in addition to hundreds of documents with "reference material" (dictionary definitions, book chapters, papers, ...) that can be read independently from the evolving conceptual network.

Supporting evolutionary development

Knowledge, like all systems, evolves by variation (recombination and mutation) and selection (Campbell, 1974). This philosophy can be straightforwardly applied to the development of knowledge in Principia Cybernetica Web. Individual nodes can be seen as pieces of knowledge, and webs of linked nodes can be seen as knowledge systems. Recombination takes place when links are changed, so that a node which was connected to one node, is now connected to another one. Mutation happens when the content of a node is changed, or a new node is created. Creation of new nodes and links by different contributors provides a continuous source of variation. If we want to maintain or improve the quality of knowledge, we also have to apply selection. Selection can be performed “manually”, by members of the editorial board using their own judgement, or automatically, by a computer program applying certain formal selection rules. The first option is limited by subjectivity and the bounded cognitive capacities of a human being, the second is limited by the difficulty to express quality judgements in a formal way, and by the rigidity of the resulting rules. The most effective approach seems to consist in a mixed human-computer system, where intuitive judgement is complemented by computing power (Heylighen, 1991b).

A possible implementation of such an approach could be found in Thagard's (1992) ECHO program. Here, the judgement of human participants determines whether two pieces of knowledge (say, propositions or arguments in a discussion) are either coherent (confirm each other) or incoherent (contradict each other). The neural network-like computer program uses these binary coherence relations to decide which of two or more competing knowledge systems is best, in the sense that its elements are most coherent with each other and with the whole of the other knowledge. Although this program was until now only used to reconstruct some key debates in the history of science, explaining why one theory eventually replaced its rival, it would seem very promising to apply such an approach to steer an on-going discussion. Although this is not really used as yet, PCP Web's annotation function allows contributors to choose a *link type* for their annotations, distinguishing arguments that support a thesis from arguments that refute it. A complex web of such arguments and counter-arguments could be analysed by ECHO-like algorithms, attracting the attention towards the more generally coherent approaches.

Another selection criterion for knowledge, besides coherence, is simplicity. Paraphrasing Ockham's Razor: all other things being equal, the simpler a knowledge system is, the better it is (Heylighen, 1997). Key aspects of this criterion can be implemented in a relatively simple way in a knowledge web. The most straightforward aspect is the ease with which a piece of knowledge can be located within a web. This is of particular importance for distributed hypertext systems, which can contain thousands or millions of interlinked nodes, making it virtually impossible to find a particular node without a priori information. Hierarchical classification has fundamental limitations, and is poorly suited for a system developed in parallel by different contributors. Free creation of associative links can provide a much richer model, but is more likely to produce labyrinthine anarchy.

We have developed a method that allows an associative hypertext network to "self-organize" into a simpler, more meaningful, and more easily usable network (Bollen & Heylighen, 1996). The term "self-organization" is appropriate to the degree that there is no external control or editor deciding which node to link to which other node: better linking patterns emerge spontaneously. The information used to create new links is not internal to the network, though: it comes from all users collectively. In that sense one might say that the network "adapts" to its users, or "learns" from the way it is used.

Algorithms for such an adaptive web can be very simple. Every potential link is assigned a certain "fitness" or "strength". For a given node, only the links with the highest fitness are shown and thus accessible to the user. Within the node, these links are ordered by strength, so that the user will encounter the strongest link first. We have developed a number of simple learning rules, inspired by the rule of Hebb for neural networks: neurons that are activated within a short time interval become more strongly connected. Similarly, nodes of the network that are consulted by the same user within a predefined interval are linked more strongly (or become linked if they were not linked as yet), while the links between nodes that are rarely used together weaken and eventually disappear. This general approach has been successfully tested on a number of experimental networks (Bollen & Heylighen, 1996), and we are now implementing it for Principia Cybernetica Web as a whole.

Discussion

This paper has provided a review of the fundamental ideas behind the Principia Cybernetica Project, and its attempts to develop an integrated evolutionary-cybernetic world view, implemented as a conceptual network accessible via the World-Wide Web. The overall success of the project, in spite of the practical and conceptual difficulties that confront any attempt at transdisciplinary integration, appears to have vindicated its approach. Yet, it is interesting to consider the remaining unsolved problems, and in particular the reasons why some of our plans have worked out better than others. Since it is difficult to objectively estimate the success of our philosophical system (*we* believe it is wide-ranging, coherent, and gives satisfactory answers to many essential questions, but people with a fundamentally different world view, e.g. mechanicists or spiritualists, are likely to disagree with that), the following discussion will focus on the practical development.

One (unplanned) reason for the success of the project was the excellent working relation and shared vision between the three members of the editorial board. In such a complex and long-term collaborative enterprise it is often difficult to avoid conflicts because of different points of view or personal incompatibilities. Yet, in spite of the differences in cultural and scientific backgrounds between the members of the PCP board, they have complemented each other remarkably well. A possibly negative effect of this close relationship is that the board has never been extended beyond the three people that were there from the beginning. Although we regularly have come into contact with other scientists, who shared an important part of our vision, their ambitions were never so close to ours that we considered inviting them to join the board, and help us carry the full responsibility for the development of PCP. Instead, we created a board of "associates", who support and advise the project, but whose major activities lie elsewhere.

More generally, the number of "non-discussion" nodes authored by contributors outside the editorial board is smaller than we had hoped for. Most contributors appear to be interested in one or a few components of the Principia Cybernetica world view, and tend to limit their contribution to those topics. An additional problem may be that these domain specialists have little incentive for publishing their ideas in Principia Cybernetica Web, since this is not a traditional journal publication for which they get scientific credit. The blurring of boundaries between journals and electronic publications may reduce this problem in the future.

On the other hand, most people who are interested in an integrated world view either do not have very concrete ideas about how it should look, or have developed their own idiosyncratic philosophies, often with little connection to what exists in relevant domains of science and philosophy. The latter approaches often come from amateurs, working outside traditional academic circles, who may have a number of interesting ideas, but lacking the culture of discussion, criticism and scepticism that defines scholarly research, tend to overestimate the value of their own contributions. This leads to a high "crackpot" factor, which makes most of these proposals useless for incorporation into PCP web. In our experience, even highly trained scientists tend to become amateurs when they leave their specialized domain and try to address wide-ranging philosophical problems. This is a general difficulty besetting all interdisciplinary research.

The most useful result of the constant shower of proposals and reactions that we receive at the Principia Cybernetica board is that it attracts our attention to various relevant developments that we were not aware of. Even if the proposal itself may not be worth following up, it may contain some pointers to related material that provides inspiration for extending the conceptual network. Also, if a number of independent proposals carry the same misunderstanding about one of PCP's fundamental concepts, then this stimulates us to incorporate a text that explicitly corrects these erroneous assumptions. The public nature of our website, which can be freely consulted by anyone interested in its subjects, allows us to reach a much wider and more diverse audience than traditional scientific journals, and this makes for a much more diverse input of suggestions and comments.

Another unanticipated reason for the success of PCP is the emergence of the World-Wide Web around 1993. Our original proposal (Joslyn, 1991) had anticipated the need for such a distributed hypermedia system to efficiently implement the project, but our expectation was that either we would have to develop a system of our own, or await the appearance of commercial solutions, either of which would have led to a much slower development than the explosive growth that we experienced with the web.

Again, this unexpectedly positive development may have had a negative side effect: the initial proposal (Joslyn, 1991) emphasized the need for node types and link types that would turn the hypermedia network into a true semantic network, allowing automatic inference. Although the creator of the web, Tim Berners-Lee, originally foresaw the use of typed links, this never developed into an accepted standard, mostly because the rapid commercial exploitation of the web led to a focus on visual appearance and other special effects meant to impress the casual visitor, rather than on improving the organization and content of the underlying knowledge network. It is to be hoped that the newly developed XML standard for the web will change this situation, and lead to the development of tools that allow the efficient manipulation of knowledge structures.

Still, we did not really need specific standards or tools to assign types to nodes and links. Apart from a small-scale experiment with a semantic network defining basic systems concepts (http://pcp.vub.ac.be/SYSCONC.html), we have not really used this opportunity. This may be in part because we underestimated the complexity of organizing knowledge into a true conceptual network with formally defined semantic categories. Yet, we have explored formalisms that allow this to be done (Heylighen, 2000), and the present popularity of research into "formal ontologies" of fundamental classes of concepts seems to indicate that the necessary formal and practical tools may be available quite soon. This would allow us to turn the Principia Cybernetica world view from an intuitively coherent philosophy into a logically or formally coherent network of concepts that could be processed by automatic programs to draw inferences, answer queries, and suggest new concepts and hypotheses (Heylighen, 2000). It is clear to us, though, that even with such powerful new tools, this will require many more years of in-depth analysis, reflection and discussion of the existing ideas.

Acknowledgments:

Although this paper basically reflects my personal view of the Principia Cybernetica Project, the results described in this paper are of course largely due to the other Principia Cybernetica contributors, and in particular my fellow members of the editorial board, Valentin Turchin and Cliff Joslyn, and our assistant editor, Johan Bollen, who implemented an important part of the system.

References

Aerts, D., Apostel L., De Moor B., Hellemans S., Maex E., Van Belle H., Van Der Veken J. (1994): Worldviews: From Fragmentation to Integration, (VUB Press, Brussels).

Ashby, Ross (1956): *Introduction to Cybernetics,* (Methuen, London)

Bollen J. & Heylighen F. (1996) “Algorithms for the Self-organisation of Distributed, Multi-user Networks. Possible application for the future World Wide Web”, in: *Cybernetics and Systems '96* R. Trappl (ed.), (Austrian Society for Cybernetics), p. 911-916.

Boulding K. (1956): "General Systems Theory - The Skeleton of Science", *General Systems Yearbook* 1, p. 11-17.

Campbell D.T. (1974a): "Evolutionary Epistemology", in: *The Philosophy of Karl Popper,* Schilpp P.A. (ed.), (Open Court Publish., La Salle, Ill.), p. 413-463.

Campbell D.T. (1974b): "'Downward causation' in Hierarchically Organized Biological Systems", in: Studies in the Philosophy of Biology, F.J. Ayala & T. Dobzhansky (ed.), (Macmillan Press), p. 179-186

Campbell D.T. (1979): "Comments on the sociobiology of ethics and moralizing", Behavioral Science 24, p. 37-45.

Dawkins R. (1976): The Selfish Gene, (Oxford University Press, New York).

Heylighen F. & Bernheim J. (2000a): "Global Progress I: empirical evidence for increasing quality of life", *Journal of Happiness Studies 1* [in press]

Heylighen F. & Bernheim J. (2000b): "Global Progress II: evolutionary mechanisms and their side-effects", *Journal of Happiness Studies 1* [in press]

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. & Campbell D.T. (1995): “Selection of Organization at the Social Level: obstacles and facilitators of metasystem transitions”, *World Futures: the Journal of General Evolution 45,* p. 181-212.

Heylighen F. & Joslyn C. (1993): “Electronic Networking for Philosophical Development in the Principia Cybernetica Project”, *Informatica* 17, No. 3, p. 285-293.

Heylighen F. (1988): "Formulating the Problem of Problem-Formulation", in: *Cybernetics and Systems '88*, Trappl R. (ed.), (Kluwer Academic Publishers, Dordrecht), p. 949-957.

Heylighen F. (1990): "A Structural Language for the Foundations of Physics", *International Journal of General Systems* 18, p. 93-112.

Heylighen F. (1991a): "Modelling Emergence", *World Futures: the Journal of General Evolution* 31 (Special Issue on *Emergence*, edited by G. Kampis), p. 89-104.

Heylighen F. (1991b): “Design of a Hypermedia Interface Translating between Associative and Formal Representations”, *International Journal of Man-Machine Studies* 35, p. 491-515.

Heylighen F. (1992a): "A Cognitive-Systemic Reconstruction of Maslow's Theory of Self-Actualization", *Behavioral Science* 37, p. 39-58.

Heylighen F. (1992b): "Principles of Systems and Cybernetics: an evolutionary perspective", in: *Cybernetics and Systems '92*, R. Trappl (ed.), (World Science, Singapore), p. 3-10.

Heylighen F. (1995): "(Meta)systems as constraints on variation", *World Futures: the Journal of General Evolution* .45, p. 59-85.

Heylighen F. (1997): "Objective, subjective and intersubjective selectors of knowledge", *Evolution and Cognition* 3:1, p. 63-67.

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. (2000): "Bootstrapping knowledge representations: from entailment meshes via semantic nets to learning webs”, Kybernetes (in press).

Heylighen F., Joslyn C. & Turchin V. (1991) : “A Short Introduction to the Principia Cybernetica Project”, *Journal of Ideas* 2, #1 p. 26-29.

Heylighen F., Joslyn C. & Turchin V. (1995) (eds.): *The Quantum of Evolution.*(Gordon and Breach Publishers, New York) (volume 45, numbers 1-4 (special issue) of "World Futures: the journal of general evolution"), 244 p.

Joslyn C. & Heylighen F.: (1999) "Cybernetics", in: *The Encyclopedia of Computer Science*, ed. J. Hemmendinger and A. Ralston, (MacMillan Reference), pp. 372-375.

Joslyn C. (1991): “Tools for the Development of Consensually-Based Philosophical Systems: a feasibility study for the Principia Cybernetica Project” (Principia Cybernetica Technical Report)

Joslyn C., Heylighen F. & Turchin V. (1993): “Synopsys of the Principia Cybernetica Project”, in: *Proc. 13th Int. Congress on Cybernetics* (Association Internationale de Cybernetique, Namur), p. 509.

Joslyn C.: (2000a) "Levels of Control and Closure in Complex Semiotic Systems", Annals of the New York Academy of Sciences, special issue on "Closure", ed. J. Chandler, G. van de Vijver, v. 901, pp. 67-74, .

Joslyn C.: (2000b) "The Semiotics of Control and Modeling Relations in Complex Systems", submitted to *Biosystems*, originally appared in: *Proc. Workshop on Control Mechanisms for Complex Systems*, ed. Michael Coombs, (NMSU Press, Las Cruces, 1997).

Kauffman S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press, New York,

Löfgren L. (1990): “On the Partiality of Self-Reference”, in: Self-Steering and Cognition in Complex Systems, Heylighen F.. Rosseel E. & Demeyere F. (eds.), (Gordon and Breach, New York), p. 47-64.

Maslow A. (1970): Motivation and Personality (2nd ed.), (Harper & Row, New York).

Maynard Smith J. & Szathmáry E. (1995): The Major Transitions in Evolution, (W.H. Freeman, Oxford)

Powers W. T. (1973): Behavior: the Control of Perception, (Aldine, Chicago).

Simon H.A. (1962): "The Architecture of Complexity", *Proceedings of the American Philosophical Society* 106, p. 467-482.

Stewart, John E. (2000): Evolution's Arrow: The direction of evolution and the future of humanity (Chapman Press, Rivett, Canberra, Australia)

Teilhard De Chardin P. (1959): The Phenomenon of Man, (Harper & Row, New York).

Thagard P. (1992): *Conceptual Revolutions*, (Princeton University Press, Princeton NJ).

Turchin V. (1977): The Phenomenon of Science, (Columbia University Press, New York).

Turchin V. (1987): "Constructive Interpretation of Full Set Theory", J. of Symbolic Logic 52:1 , p. 172-201.

Turchin V. (1991): "Cybernetics and Philosophy", in: *The Cybernetics of Complex Systems*, F. Geyer (ed.), (Intersystems, Salinas, California), p. 61-74.

Turchin V. (1993a): "On Cybernetic Epistemology", Systems Research 10:1, p. 3-28.

Turchin V. (1993b): "The Cybernetic Ontology of Actions", Kybernetes 22:2, p. 10-30.

Van Gigh J.P. (ed.) (1986): *Decision-making about Decision-making: metamodels and metasystems*, (Abacus Press, Cambridge).

Varela, F., Principles of Biological Autonomy, North Holland, New York, 1979.

von Bertalanfy L. (1968): *General Systems Theory*, (Braziller, New York).

von Foerster H. (1979): The Cybernetics of Cybernetics (Gordon and Breach, New York).

Whitehead A.N. & Russell B. (1910-1913): *Principia Mathematica* (vol. 1-3), (Cambridge University Press, Cambridge).

Whitehead A.N. (1929): Process and Reality: an essay in cosmology, (Cambridge University Press, Cambridge).

Biography:

Dr. Francis Heylighen received his university degree in mathematical physics in 1982, and his Ph.D. in 1987 from the Free University of Brussels (VUB). He is a Research Associate for the Fund for Scientific Research - Flanders (FWO) and a co-director of the Center "Leo Apostel" at the VUB. His research focused first on the foundations of physics, then on the evolution of knowledge and of complexity, which he addresses from a cybernetic perspective. Together with his collaborator Johan Bollen, he has developed a self-organizing, "learning" website, implementing some of these ideas. Most recently, he has extended the underlying principles to understand the future evolution of the information society. Dr. Heylighen has authored some 80 scientific publications. He is editor of the Principia Cybernetica Project, chairman of the Global Brain Group, and member of the editorial boards of the "Journal of Memetics", "Informatica" and "Entropy".