
Principles of Systems and cybernetics:
an evolutionary perspective

Francis HEYLIGHEN*

PO, Free University of Brussels, Pleinlaan 2, B-1050 Brussels,
Belgium; fheyligh@vnet3.vub.ac.be

ABSTRACT: A set of fundamental principles for the cybernetics domain is sketched, based on the spontaneous emergence of systems through variation and selection. The (mostly self-evident) principles are: selective retention, autocatalytic growth, asymmetric transitions, blind variation, recursive systems construction, selective variety, requisite knowledge and incomplete knowledge. Existing systems principles, such as self-organization, "the whole is more than the sum of its parts", and order from noise can be reduced to implications of these more primitive laws. Others, such as the law of requisite variety, the 2nd law of thermodynamics, and the law of maximum entropy production are clarified, or restricted in their scope.

Introduction

Principles or laws play the role of expressing the most basic ideas in a science, establishing a framework or methodology for problem-solving. The domain of General Systems and Cybernetics is in particular need of such principles, since it purports to guide thought in general, not just in a specific discipline. Unfortunately, the few generally used principles of the domain, such as the law of requisite variety, or the principle that the whole is more than the sum of its parts, are typically ambiguous or controversial, and lack coherence with each other.

The present work purports to start a general examination of principles of cybernetics and systems, within the framework of the Principia Cybernetica Project (Heylighen, Joslyn & Turchin, 1991; Turchin, 1991). The Principia Cybernetica philosophy is evolutionary: systems and their cybernetical organization are constructed through the self-organizing process of blind variation and natural selection. This process function as a skeleton interconnecting all principles.

The study will on the one hand critically assess existing principles, clarifying their meaning, on the other hand try to formulate new principles which may generalize or interconnect known laws. The ultimate goal is to arrive at a network of concepts and principles similar to a formal system, with "axioms" implicitly defining primitive concepts, definitions of higher order concepts, and "theorems", derived from the more primitive axioms and definitions. The fundamental principles, like all good axioms, are supposed to be self-evident, if not tautologous. Their implications, like most theorems, on the other hand, may be far from trivial, and sometimes even counter-intuitive.

This paper will propose a first, necessarily limited and sketchy, overview of the principles that I think are most basic, starting from the most primitive ones, and building up towards less obvious ones. This overview is offered for discussion and elaboration by other systems researchers. A more in-depth treatment of this issue is being prepared in the form of a series of journal papers (Heylighen, forthcoming).

The Principle of Selective Retention

Stable configurations are retained, unstable ones are eliminated.

This first principle is tautological in the sense that stability can be defined as that what does not (easily) change or disappear. Instability then is, by negation, that what tends to vanish or to be

replaced by some other configuration, stable or unstable. The word "configuration" denotes any phenomenon that can be distinguished. It includes everything that is called feature, property, state, pattern, structure or system.

The principle can be interpreted as stating a basic distinction between stable configurations and configurations undergoing variation. This distinction has a role in evolution which is as fundamental as that between A and not A in logic. Without negation, we cannot have a system of logic. Without (in)stability we cannot describe evolution. The tautology plays a role similar to the principle of contradiction: "A and not A cannot both be true". The distinction between stable and changing is not as absolute as that between A and not A, though. We do not require a principle of the excluded middle, since it is clear that most configurations are neither absolutely stable nor absolutely unstable, but more or less stable. In this more general formulation, the principle would read:

More stable configurations are less easily eliminated than less stable ones

The Principle of Autocatalytic Growth

Stable configurations that facilitate the appearance of configurations similar to themselves will become more numerous

This self-evident principle is the companion of the principle of selective retention. Whereas the latter expresses the conservative aspect of evolution, maintenance or survival, the former expresses the progressive aspect, growth and development. Autocatalytic growth describes as well biological reproduction, as the positive feedback or non-linearity characterizing most inorganic processes of self-organization, such as crystal growth. The principle simply states that it suffices for a configuration to be stable, and in some respect autocatalytic or self-replicating, in order to undergo a potentially explosive growth.

Such configurations, in biology, are said to have a high fitness and that gives them a selective advantage over configurations with a lower fitness. The fact that growth requires (finite) resources implies that growth must eventually stop, and that two configurations using the same resources will come in competition for these resources. Normally the fitter configuration will outcompete the less fit one, so that no resources are left for the latter (survival of the fittest). Such a generalization of the principle of selective retention may be called the principle of natural selection.

The Principle of Asymmetric Transitions: entropy and energy

A transition from an unstable configuration to a stable one is possible, but the converse is not.

This principle implies a fundamental asymmetry in evolution: one direction of change (from unstable to stable) is more likely than the opposite direction. The generalized, "continuous" version of the principle is the following:

The probability of transition from a less stable configuration A to a more stable one B is larger than the probability for the inverse transition: $P(A \rightarrow B) > P(B \rightarrow A)$ (under the condition $P(A \rightarrow B) \neq 0$)

A similar principle was proposed by Ashby in his Principles of the Self-Organizing System (1962): "We start with the fact that systems in general go to equilibrium. Now most of a system's states are non-equilibriumal [...] So in going from any state to one of the equilibria, the system is going from a larger number of states to a smaller. In this way, it is performing a selection, in the purely objective sense that it rejects some states, by leaving them, and retains some other state, by sticking to it. "

This reduction in the number of reachable states signifies that the variety, and hence the statistical entropy, of the system diminishes. It is because of this increase in negentropy or organization that Ashby calls the process self-organization. But how does this fit in with the 2nd law of thermodynamics, which states that entropy in closed systems cannot decrease? The easy way out is to conclude that such a self-organizing system cannot be closed, and must lose entropy to its environment (von Foerster, 1960).

A deeper understanding can be reached by going back from the statistical definition of entropy to the thermodynamic one, in terms of energy or heat. Energy is defined as the capacity to do work, and working means making changes, that is to say exerting variation. Hence energy can be viewed as potential variation. A stable configuration does not undergo variation. In order to destroy a stable equilibrium, you need to add energy, and the more stable the configuration, the more energy you will need. Therefore stability is traditionally equated with minimal energy.

The 1st law of thermodynamics states that energy is conserved. A naive interpretation of that law would conclude that the principle of asymmetric transitions cannot be valid, since it postulates a transition from an unstable (high energy) to a stable (low energy) configuration. If energy is absolutely conserved, then an unstable configuration can only be followed by another unstable configuration. This is the picture used in classical mechanics, where evolution is reversible, that is to say symmetric. Incidentally, this shows that the principle of asymmetric transitions is not tautological - though it may appear self-evident - , since a perfectly consistent theory (classical mechanics) can be built on its negation.

Thermodynamics has enlarged that picture by allowing energy dissipation. But what happens with the "dissipated" energy? A simple model is provided by a quantum system (e.g. an electron bound in an atom) with its set of - usually discrete - energy levels. A configuration at a higher level will spontaneously fall down to a lower level, emitting a photon which carries the surplus energy away. In order to bring back the electron to its higher level, energy must be added by having a photon of the right energy and direction hit the electron, a rather improbable event. Hence, the low level can be viewed as a stable configuration, with a small probability of transition.

The conjunction of energy conservation and asymmetric transitions implies that configurations will tend to dissipate energy (or heat) in order to move to a more stable state. For a closed system, this is equivalent to the thermodynamical interpretation of the 2nd law, but not to the statistical one, as the statistical entropy can decrease when transition probabilities are asymmetric. In an open system, on the other hand, where new energy is continuously added, the configuration will not be able to reach the minimum energy level. In that case we might assume that it will merely tend to maximally dissipate incoming energy, since transitions where energy is emitted are (much) more probable than transitions where energy is absorbed. That hypothesis seems equivalent to the Law of maximum entropy production (Swenson, 19), which describes dissipative structures and other far-from-equilibrium configurations. In such configurations the stability is dynamic, in the sense that what is maintained is not a static state but an invariant process.

Such an application of the principle of asymmetric transitions is opposite to the most common interpretation of the 2nd law, namely that disorder and with it homogeneity tend to increase. In the present view, configurations tend to become more and more stable, emitting energy in the process. This might be seen as a growing differentiation between the negative energy of stable bonds, and the positive energy of photons and movement. Recent cosmological theories hypothesize a similar spontaneous separation of negative and positive energies to account for the creation of the universe

out of a zero-energy vacuum (Hawking, 1988).

The Principle of Blind Variation

At the most fundamental level variation processes "do not know" which of the variants they produce will turn out to be selected

This principle is not self-evident, but can be motivated by Ockham's razor. If it were not valid, we would have to introduce some explanation (e.g. design by God) to account for the "foreknowledge" of variation, and that would make the model more complicated than it needs to be. The blindness of variation is obvious in biological evolution, based on random mutations and recombinations. Yet even perfectly deterministic dynamical systems can be called blind, in the sense that if the system is complex enough it is impossible to predict whether the system will reach a particular attractor (select a stable configuration of states) without explicitly tracing its sequence of state transitions (variation) (Heylighen, 1991).

Of course many interactions are not blind. If I tackle a practical problem, I normally do not try out things at random, but rather have some expectations of what will work and what will not. Yet this knowledge itself was the result of previous trial-and-error processes, where the experience of success and failure was selectively retained in my memory, available for guiding later activities. Similarly, all knowledge can be reduced to inductive achievements based on blind-variation-and-selective-retention (BVSR) at an earlier stage. Together with Campbell (1974), I postulate that it must be possible to explain all cases of "non-blindness" (that is to say variation constrained in such a way as to make it more likely to satisfy selection) as the result of previous BVSR processes.

The BVSR formula summarizes three previous principles: selective retention, asymmetric transitions, and blind variation. The second principle is implicit in the "and" of "blind-variation-and-selective-retention", since it ensures that configurations produced by blind variation can make the transition to selective retention, unlike configurations in classical mechanics which remain unstable.

The Principle of Selective Variety

The larger the variety of configurations a system undergoes, the larger the probability that at least one of these configurations will be selectively retained.

Although this principle is again self-evident or tautologous, it leads to a number of useful and far from trivial conclusions. For example, the less numerous or the farther apart potential stable configurations are, the more variation (passing through a variety of configurations) the system will have to undergo in order to maintain its chances to find a stable configuration. In cases where selection criteria, determining which configurations are stable and which are not, can change, it is better to dispose of a large variety of possible configurations. If under a new selective regime configurations lose their stability, a large initial variety will make it probable that at least some configurations will retain their stability. A classic example is the danger of monoculture with genetically similar or identical plants: a single disease or parasite invasion can be sufficient to destroy all crops. If there is variety, on the other hand, there will always be some crops that survive the invasion.

Another special case is the "order from noise" principle (von Foerster, 1960), related to "order out of chaos". Noise or chaos can here be interpreted as rapid and blind variation. The principle states that addition of such noise makes it more likely for a system to evolve to an ordered (stable) configuration. A practical application is the technique of (simulated) annealing, where noise or variation is applied in stepwise decreasing amounts, in order to reach a maximally stable configuration.

BVSR processes recursively construct stable systems by the recombination of stable building blocks

The stable configurations resulting from BVSR processes can be seen as primitive elements: their stability distinguishes them from their variable background, and this distinction, defining a "boundary", is itself stable. The relations between these elements, extending outside the boundaries, will initially still undergo variation. A change of these relations can be interpreted as a recombination of the elements. Of all the different combinations of elements, some will be more stable, and hence will be selectively retained.

Such a higher-order configuration might now be called a system. The lower-level elements in this process play the role of building blocks: their stability provides the firmness needed to support the construction, while their variable connections allow several configurations to be tried out. The principle of "the whole is more than the sum of its parts" is implied by this systemic construction principle, since the system in the present conception is more than a mere configuration of parts, it is a stable configuration, and this entails a number of emergent constraints and properties (Heylighen, 1991). A stable system can now again function as a building block, and combine with other building blocks to form an assembly of an even higher order, in a recursive way.

Simon (1962) has argued in his famous "The Architecture of Complexity" that such stable assemblies will tend to contain a relatively small number of building blocks, since the larger a specific assembly, the less probable that it would arise through blind variation. This leads to a hierarchical architecture, that can be represented by a tree.

Two extensions must be made to the Simon argument (cf. Heylighen, 1989). 1) If one takes into account autocatalytic growth, as when a small stable assembly makes it easier for other building blocks to join the assembly, the number of building blocks at a given level can become unlimited. 2) It is possible, though less probable, that a given building block would participate in several, overlapping stable assemblies; it suffices that its configuration would satisfy two (or more) selection criteria, determining stable systems. It is clear, however, that the more selection criteria a configuration would have to satisfy, the less likely that such a configuration would be discovered by blind variation. These two points lead us to generalize the tree structure of Simon's "nearly-decomposable" architecture to a loose or quasi-hierarchy (Joslyn, 1991), which in parts can be very flat, and where some nodes might have more than one mother node.

Control systems

The previous principles provide a set of mechanisms describing the spontaneous emergence and self-organization of multilevel systems, becoming ever more stable (in a generalized, 'dynamical' sense), more fit, and more complex. Control systems are a specific type of such multilevel systems, where a stable configuration is maintained by selectively counteracting perturbations. There is no space here to examine in detail how control systems emerge through BVSR, but the issue can be clarified by considering the concept of an anticipatory or vicarious selector (Campbell, 1974).

A selector is a stable system capable of selecting variation. A vicarious selector carries this selection out in anticipation of something else, e.g. the environment or "Nature" at large. For example, molecule configurations selectively retained by a crystal template are intrinsically stable, and would have been selected even without the presence of a template. The template basically accelerates (catalyses) selection, and thus can be said to anticipate, or to vicariously represent, the naturally selected configuration. The selection made by a template is invariant.

However, one can also imagine anticipatory selectors making different selections under different circumstances, compensating different perturbations by different actions. This anticipatory selection has the advantage that inadequate internal variations will no longer lead to the destruction of the system, since they will be eliminated before the system as a whole becomes unstable.

This mechanism can be illustrated by considering what Powers (1989) calls the most primitive example of a control system, a bacterium that changes the rate of random variation of its movements in function of the favourableness of its environment. When the concentration of food increases, its variation of movement becomes small. When the concentration of food decreases (or that of poison increases), there is a strong variation. The only selection the bacterium makes is that between moving in the same direction (selective retention), or changing course (blind variation). That selection anticipates the natural selection that would happen if the bacterium was passive (that is to say, if it was not exerting control): if it would stay long enough in the unfavourable place, it would die; if it would move to a more favourable place it would survive. The bacterium is in fact applying the principle of selective variety: it increases variation when the chances of being selectively retained become less. This internally directed, selective counteraction of perturbations from a stable configuration can be taken as a definition of control. This leads us straightforwardly to a derivation of some classic principles of control.

The Law of Requisite Variety

The larger the variety of actions available to a control system, the larger the variety of perturbations it is able to compensate.

This is another application of the principle of selective variety, formulated above. However, a stronger form of Ashby's Law (1958), "the variety in the control system must be equal to or larger than the variety of the perturbations in order to maintain stability", does not hold in general. Indeed the underlying "only variety can destroy variety" assumption is in contradiction with the principle of asymmetric transitions which implies that spontaneous decrease of variety is possible. For example, the bacterium described above disposes of a minimal variety of only two actions: increase or decrease the rate of random movements. Yet, it is capable to cope with a quite complex environment, with many different types of perturbations (Powers, 1989). Its blind "transitions" are normally sufficient to find a favourable ("stable") situation, thus escaping all dangers.

The Law of Requisite Knowledge

In order to adequately compensate perturbations, a control system must "know" which action to select from the variety of available actions.

This principle reminds us that a variety of actions is not sufficient for effective control, the system must be able to (vicariously) select an appropriate one. Without knowledge, the system would have to try out an action blindly, and the larger the variety of perturbations, the smaller the probability that this action would turn out to be adequate. Notice the tension between this law and the previous one: the more variety, the more difficult the selection to be made, and the more complex the requisite knowledge. "Knowing" signifies that the internal (vicarious) selector must be a model or representation of the external, potentially selecting perturbations. Ideally, to every class of perturbations there corresponds a class of adequate counteractions. This correspondence might be represented as a homomorphism from the set of perturbations to the set of (equivalence classes of) compensations. In the case of the bacterium, the class of favourable situations is mapped onto the action "decrease variation", whereas unfavourable situations are mapped onto "increase variation".

However, this does not imply that knowledge would consist of a homomorphic image of the objects in the environment. Only the (perturbing) processes of the environment need to be represented, not its static structure.

An equivalent principle was formulated by Conant and Ashby (1970) as "Every good regulator of a system must be a model of that system". Therefore the present principle can also be called the law of regulating models.

The Principle of Incomplete Knowledge

The model embodied in a control system is necessarily incomplete

This principle can be deduced from a lot of other, more specific principles: Heisenberg's uncertainty principle, implying that the information a control system can get is necessarily incomplete; the relativistic principle of the finiteness of the speed of light, implying that the moment information arrives, it is already obsolete to some extent; the principle of bounded rationality (Simon, 1957), stating that a decision-maker in a real-world situation will never

have all information necessary for making an optimal decision; the principle of the partiality of self-reference (Loefgren, 1990), a generalization of Goedel's incompleteness theorem, implying that a system cannot represent itself completely, and hence cannot have complete knowledge of how its own actions may feed back into the perturbations. As a more general argument, one might note that models must be simpler than the phenomena they are supposed to model. Otherwise, variation and selection processes would take as much time in the model as in the real world, and no anticipation would be possible, precluding any control. Finally, models are constructed by blind variation processes, and, hence, cannot be expected to reach any form of complete representation of an infinitely complex environment.

Acknowledgments: I thank C. Joslyn, V. Turchin and other Principia Cybernetica contributors for a preliminary discussion, inciting me to clarify many points in the draft.

References

Ashby W.R. (1958): "Requisite Variety and Implications for Control of Complex Systems", *Cybernetica* 1, p. 83-99.

Ashby W.R. (1962): "Principles of the Self-Organizing System", in: *Principles of Self-Organization*, von Foerster H. & Zopf G.(eds.), (Pergamon, Oxford), p. 255-278.

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

Conant R.C. and Ashby W.R. (1970): "Every Good Regulator of a System Must Be Model of that System", *Int. J. Systems Science* 1:2, p. 89-97.

Hawking S.W. (1988): *A Brief History of Time*, (Bantam, London).

Heylighen F. (1989): "Self-Organization, Emergence and the Architecture of Complexity", in: *Proc. 1st Eur. Conf. on System Science*, (AFCET, Paris), p. 23-32.

Heylighen F. (1991): "Modelling Emergence", *World Futures* 31, p. 89-104.

Heylighen F. (forthcoming): "Principles of Evolution and Self-organization", to be submitted to *Int. Journal of General Systems*.

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

Joslyn C. (1991): "Hierarchy and Strict Hierarchy in General Information Theory", in: Proc. 1991 Congress of the International Society for Systems Science.

Loefgren L. (1990): "On the Partiality of Self-Reference", in: Self-Steering and Cognition in Complex Systems, Heylighen et al. (eds.), (Gordon & Breach, NY), p.47-64.

Powers, W.T. (1989): "An Outline of Control Theory", in: Living control systems, (Control Systems Group, Gravel Switch: KY), p. 253-293.

Simon H.A. (1957): Models of Man : Social and Rational, (Wiley, London).

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

Swenson R. (1989): "Emergent Attractors and the Law of Maximum Entropy Production: foundations to a theory of general evolution", Systems Research 6:3, p. 187-198.

Turchin V. (1990): "Cybernetics and Philosophy", in: The Cybernetics of Complex Systems, F. Geyer (ed.), (Intersystems, Salinas, California), p. 61-74.

von Foerster H. (1960): "On Self-Organizing Systems and their Environments", in: Self-Organizing Systems, Yovitts M.& Cameron S. (ed.), (Pergamon, New York), p.31-50.

* Senior Research Assistant NFWO (Belgian National Fund for Scientific Research)