c-institute

Paper/Article name:	Evolutionary fitness for agents and agencies
Date of Publication:	April 2006
Overview:	The paper explores the notion of evolutionary fitness as applied cognition and responses

Evolutionary Fitness for agents and agencies

The extension of Darwin's evolutionary thinking to various other strands of science (e.g., evolutionary economics, Saviotti & Metcalfe, 1992; evolutionary psychology, Buss, 1991; evolutionary computation, and the evolution of chemical compounds or elementary particles), is now quite common. One of the first was evolutionary epistemology, defined by Campbell (1974), citing that all knowledge is the product of variation and natural selection. Applicable to the primitive knowledge stored in the genes (allowing organisms to adapt to its environment), and more sophisticated theories of science, which undergo variation when a scientist speculates and selection when inadequate theories fail to pass empirical tests. The Principia Cybernetica Project (Heylighen, Joslyn & Turchin, 1991) aims to carry the study of the evolutionary origin of systems to its logical end point, extending it beyond, cognitive, social, biological, and include interconnections between systems of different types. This should lead to reconstruct the extended chain of variation and selection processes producing complexity, from elementary particles to the intricate structures of society. This approach has been called "Metasystem Transition Theory" (Heylighen, Joslyn & Turchin, 1991).

Heylighen (1992), applies this philosophy to the origin of knowledge, studying the evolutionary preconditions that make knowledge at all possible. He raises a fundamental question every epistemologist should ask, i.e., given that knowledge consists of extremely simple models of an infinitely complex reality, how can we explain that knowledge is still most of the time reliable? He proceeds to answer the question by

linking the mechanism of default reasoning to the natural selection of cognised phenomena. Evolutionary epistemology (Campbell, 1974) assumes that the function of knowledge is to secure the survival or "fitness" of an organism. Further to this, cybernetics adds that survival in a variable environment is a control problem, which requires adequate compensation of perturbations that make the system deviate from its goal (maintaining or increasing fitness). Perhaps the most famous principle of control is Ashby's Law of Requisite Variety (1958), which states that, in order to achieve complete control, the variety of compensatory actions a control system is capable to execute must be at least as great as the variety of perturbations that might occur. Adding Conant and Ashby's (1970) principle that "every good regulator (controller) of a system must be a model of that system", together, imply that perfect control can only be established if there exists a one-to-one (bijective) mapping from the set of all possible perturbations to the set of all counteractions the system can execute (Heylighen, 1992).

The question this view raises then is: how can an infinitely extended world be isomorphically mapped onto a finite (our mind) cognitive system? The traditional answer is that the concrete mapping is not isomorphic (one-to-one) but homomorphic (many-to-one): most details or distinctions about the world are filtered out (e.g., two slightly different frequencies of light will both be perceived as "red" and stored under that single heading in a person's brain). Under Conant and Ashby's analysis, many-to-one mappings, imply loss of information, equals loss of control (the less information survives the mapping, the greater the variation in the outcome of the actions, hence a larger fluctuation around the goal state). This relation between the varieties (V) of external perturbations (E), actions (A), and outcomes (O) can be derived from Ashby's (1958) law that "only variety can destroy variety": $V(O) \geq V(E) - V(A)$.

It implies that an infinite variety of perturbations, controlled by a finite variety of actions, will still leave an infinite variety of outcomes. We must conclude that knowledge, viewed as a finite map of an infinite territory, appears like a very limited tool for achieving control. Heylighen (1992) explain why knowledge nevertheless seems so effective:

(1) Ashby's law of requisite variety should not be taken as an absolute requirement. When the system is in a naturally stable state, the variety of outcomes will be automatically decreased, even without intervening actions (citing his "principle of asymmetric transitions", Heylighen, 1992).

 (2) The present cybernetic view of knowledge does not assume a mapping between (static) objects and their representations, but between (dynamic) perturbations and actions. But what constitutes a perturbation? Apparently we only need to take into account processes that have a causal influence on the system's goal. Replacing the absolute causality of classical mechanics, where every event affects every subsequent event, by an irreversible, thermodynamic view of processes, we may conclude that most causal signals will be dissipated before they propagate to the system. In that case, what counts as "perturbations" will be a small subset of all physical events that have direct causal links to the variables defining the system's (subjective) goal.

(3) Even an infinite remaining variety of outcomes do not imply inadequate control, as long as the few "essential variables" distinguishing between survival and death are kept within bounds (e.g., variation of the organism's horizontal position have little affect upon its chances for survival, but variation in its body temperature, must be controlled within strict limits).

Heylighen (1992), argues that a small set of actions may still be sufficient to adapt to an infinitely complex environment. Every new phenomenon must somehow be put in the appropriate category, which can then be linked to an action adequate for that class of events. The function of the cognitive system is to map specific combinations of these attributes onto more abstract categories, which can then be interpreted in terms of required actions (e.g., the combination "hot", "high", "light" may elicit the concept "sun", which may trigger the action "go into the shade"). In order to adequately steer the organism towards its goal of increasing fitness, a maximum number of combinations of attributes must be put into categories denoting possible dangers (e.g. fire, predators, cliffs, rivers, insects, etc) or resources and opportunities (e.g. food, mates, water, shelters, etc). As implied by Ashby's law, the larger the variety of action-triggering categories available to the organism, the larger the control it can exercise, in succeeding in improving its fitness, and survival. Evolution thus tends to increase the number of perceivable attributes and categories.

A basic mechanism for minimising the complexity of categorisation is default reasoning. Each time a new combination of attributes is encountered, the organism must find the appropriate category in which to fit the perceived phenomenon, in order to further infer appropriate actions. The same combination might fit several categories, or none at al. Rather than systematically test all categories (Is it a bird? Is it a plane? Is it...?), the organism will immediately pick up the "most likely" category, until it encounters evidence that another categorisation is needed. It will then try out the "second most likely" category, and so on.

The classical example of default reasoning is the assumption that if something is a bird (earlier experience or attribute), it can be expected to fly (inferred category). This is probably true in over 99% of cases; yet, the existence of ostriches and penguins clearly shows that this is not a universal truth. Violations of the default expectation will be encoded in the cognitive system as exceptions: if it is a bird, then it can fly, except if it is penguin or an ostrich. But the awareness of the "exceptional" situation will trigger new default expectations: if it is an ostrich, then it can run; or, if it is a penguin, then it can swim. Again, there will be exceptions to these rules: if it is an ostrich, and it has a broken leg, it cannot run. But that expectation might in very unusual circumstances again be violated: perhaps an ostrich with a broken leg could still run if it was wearing some kind of artificial support?

The system behind this type of reasoning is called a default hierarchy (Holland et al., 1986), and consists of different levels of expectations, with the most likely being at the top, the less likely below. As ones goes deeper down into the exceptions and exceptions to exceptions, more attributes are added to the necessary conditions, and thus triggering conditions become more specific, and less probable (e.g., it is unlikely to encounter a penguin wearing a fin brace?).

Complexity of decision-making

The complexity of decision-making or categorisation (the average number of alternate categories to be explored before the appropriate one is found), appear to be a good measure for cognitive effort, and is similar to the view used by Simon (1962) in his paper on the "Architecture of Complexity" when arguing that hierarchical decomposition enormously decreases the complexity of problem-solving (Heylighen, 1991b).

c-institute

Cognitive complexity of choice between a given number of alternatives, may be kept small by:

Robust ordering: a factor that depends on the organism's capacity to learn, (storing experience as to the frequency with which a particular alternative is encountered). It is known that past frequency of occurrence, implying likeliness of future occurrence, is a fundamental determinant of learning. The "strength" of a connection determines the likeliness that the connection is later activated, and thus the (average) speed with which the alternative represented by that connection is explored.

Low entropy: a factor that partly depends on the organism, and partly on its environment. In a high entropy environment, where all types of phenomena or situations are about equally likely, cognitive complexity would be maximal, and control through knowledge would be virtually impossible. In a "mixed entropy" environment, where phenomena are equally likely, while others have strongly differentiated probabilities, cognitive complexity could be kept within bounds by filtering out all the high entropy categories.

In conclusion, the evolutionary principle of the "survival of the fittest" explains the existence of alternatives that are much more likely than others, playing the role of the "defaults". The most fundamental type of default assumption is that most likely a phenomenon is fit. For example, lame ostriches or penguins unable to swim are unfit, and will not survive (it will die because of starvation, not being able to catch fish). More generally, given the constraint that a system is a bird (having wings and feathers), we may assume that it will be most fit when it can fly, though there are exceptions, e.g., ostriches and penguins, where flight is not necessary for fitness. The "defining features" of a particular category (e.g. birds) determine a limited domain of "fit" configurations.

Most instances of the category will be kept within that domain by natural selection, and this explains the adequacy of default assumptions about members of that category. Not all configurations encountered will be fit, as variation will continuously produce fluctuations or deviations from the "normal", fit configuration. That is why default assumptions remain just that, they do not become "absolute truths", and for every rule there will always be an exception.

A physical example of the biological one above may be: "stones are hard" is a typical default assumption, i.e., for stones; hardness is a part of its fitness (soft stones will turn into sand under erosion, and will be quickly "die"). Therefore, the "abstract" default assumption that a stone is fit implies that it is hard. Another example, "water is liquid", yet its fitness depends on the environment (it is only valid when temperatures are between freezing and boiling). Another example: "people understand language", in present society being unable to understand language is unfit (e.g., a child will be taught language understanding, else it will be isolated in a home for the mentally handicapped).

References

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

Buss D.M. (1991): "Evolutionary Personality Psychology", Am. Rev. of Psychology 42, p. 459-491.

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.

Heylighen F, Joslyn C. & Turchin V. (eds) (in print): The Theory of Metasystem Transitions (special issue of "World Futures").

Heylighen F. (1991a): "Cognitive Levels of Evolution: pre-rational to meta-rational", in: The Cybernetics of Complex Systems, F. Geyer (ed.), (Intersystems, Salinas, California), p. 75-91.

Heylighen F. (1991b): "Coping with Complexity", Systemica 8, part 1, p.39-55.

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

Heylighen, F. (1992), The evolutionary basis for cognitive complexity reduction, PO, Free University of Brussels, Pleinlaan 2, B-1050 Brussels, Belgium.

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

Holland J.H., Holyoak K.J., Nisbett R.E. & Thagard P.R. (1986): Induction (MIT Pr ess, MA).

Koopman B. (1978): "Entropy Increase and Group Symmetry", in: Maximum Entropy Formalism, M. Tribus & I. Levine (eds.), (MIT Press, Cambridge, Mass.).

Saviotti P.P. & Metcalfe J.S. (eds.) (1992): Evolutionary Theories of Economic and Technological Change, (Scientific and Technical Book Service, New York).

Simon H.A. (1962): "The Architecture of Complexity", Proc. Am. Philos. Soc. 106, p. 467-482.

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