

## Decision Making in Evolving Artificial Systems Richard Walker<sup>†</sup>, Maurizio Cardaci<sup>‡</sup>

<sup>†</sup>II University of Naples

[r.walker@unina2.it](mailto:r.walker@unina2.it)

<sup>‡</sup>University of Palermo

[Cardaci@unipa.it](mailto:Cardaci@unipa.it)

### 1. INTRODUCTION

The theme of this workshop is artificial perception. In this paper we will argue that the *ecological* function of perception is to serve decision-making. If this is so the mechanisms chosen to implement perception, in natural or artificial systems, will be constrained by the requirements of decision-making and theories of decision-making will inevitably influence theories of perception. In what follows we will look at decision-making from what we hope is a new perspective, applying concepts and techniques developed by what we will call “new artificial intelligence”. We will begin, in Part 2 of the paper, with a review of traditional, “normative” theories of decision-making and of the mounting body of experimental evidence, showing the inadequacy of these theories. In Part 3 we will outline an alternative ecological/biological view of decision-making, inspired by computer models from the fields of “Artificial Neural Networks”, “Artificial Evolution” and “Evolutionary Robotics. Finally, in part 4, we will briefly examine how far this new view may be considered to be biologically realistic and examine some of the possible consequences for theories of perception and of individual perceptive systems.

### 2 TRADITIONAL THEORIES OF DECISION-MAKING

Traditional theories of decision-making have their roots in classical and neo-classical economics and are essentially *normative* in nature [Shafir & Tversky, 1995]. The goal of these theories is to describe how rational decision-makers should take their decisions and to prescribe criteria for what should, and should not be considered, as *rational* behavior.

The theory of decision-making has been traditionally formulated as a “theory of rational choice”. In this view a rational decision-maker “scores” the expected consequences of a set of alternative policies. Policies leading to favourable outcomes receive a high score and policies leading to unfavourable outcomes a low one. When all possible policies have been examined the decision-maker makes a *rational choice*, choosing the policy with the highest score.

This view of decision-making has a number of implications. In particular:

- Rational choices are based, by definition, on objective facts; they should therefore be *description invariant*. Given a set of possible policies and their known consequences two rational decision-makers, with the same preferences, should calculate the same scores and arrive at the same decision, regardless of the way in which the problem is formulated.

- Preferences should exhibit *transitivity*: if a decision-maker prefers policy A to policy B and policy B to policy C she should prefer policy A to policy C.

Classical experiments conducted during the 1970s and 1980s, have shown that these predictions are false.

In one experiment [Tversky & Kahneman, 1981] two groups of volunteers were asked to consider alternative plans to fight a viral epidemic expected to cause 600 cases of a potentially fatal disease. The first group was told that Policy A would save 200 lives whereas Policy B would either completely fail, leading to 600 deaths, or prove 100% effective, saving all 600 patients. The majority of volunteers in this group chose Policy A.

The second group of volunteers was told that with Policy A there would be 400 deaths; policy B was described in the same way as for the first group. This second group chose Policy B.

In reality the two policies are identical. With 600 cases two hundred successes imply four hundred deaths and vice-versa. The difference in the reaction of the two groups seems to depend on the different ways in which the problem is formulated. This is a clear *violation of description invariance*.

A second classic experiment tested the prediction of transitivity [Tversky 1969]. Here volunteers were asked to simulate the selection of candidates for admission to university. The selection was to be based on candidates' scores for intelligence, mental stability and social skills. The experiment was designed so that candidate B received a higher score for intelligence than candidate A, candidate C had a higher score than candidate B and so on, whereas the scores for stability and social skills showed the opposite trend with the lowest scores going to E and the highest to A.

During the experiment candidates were presented two at a time and volunteers were asked to indicate which of the two candidates they would select. In general volunteers showed a preference for high stability and good social skills: A was preferred to B, B to C, C to D and D to E. When however the volunteers were asked to choose between A and E a large majority chose E. Here again the experiment showed that real human decision-makers fail to satisfy the predictions of the theory of rational choice, in this case the prediction of transitivity.

Psychologists have presented a number of theories to explain the conflict between the predictions of normative theory and observed behavior. So-called "Prospect Theory", for instance, [Kahneman & Tversky, 1979] argues that decision-makers tend to ignore small differences (i.e. the relatively small differences in the intelligence scores of A and B, B and C etc.) making their choices on the basis of major differences (such as the large difference in the intelligence scores of A and E). A number of alternative explanations have also been offered. These explanations are theories of *imperfect rationality*. Human beings, we are asked to believe, are not clever enough or do not have enough information to achieve the higher rationality of the economist.

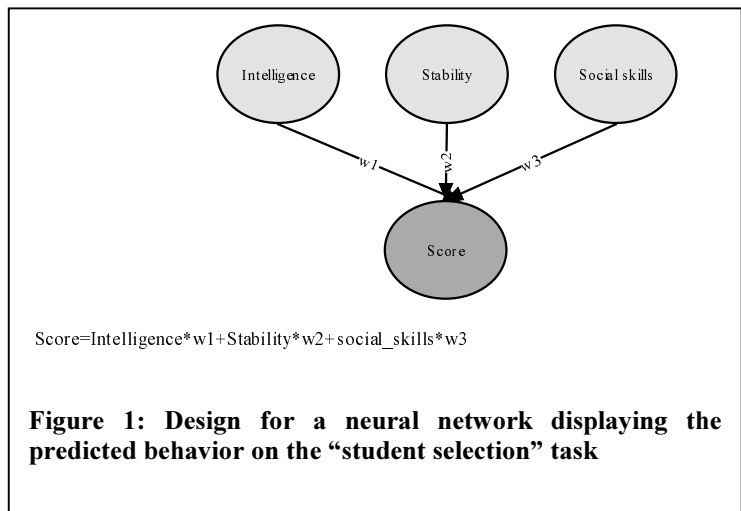
### **3 A BIOLOGICAL/ECOLOGICAL VIEW OF DECISION-MAKING: INSIGHT FROM "NEW ARTIFICIAL INTELLIGENCE"**

From a biological or an ecological viewpoint a "*rational*" *decision-making procedure* is a procedure, which, when applied repeatedly, benefits the reproductive fitness of the decision-maker. In what follows we use theoretical and experimental results from "New Artificial Intelligence" to suggest that there exist organisms whose decision-making is perfectly rational in this biological sense yet whose behavior fails to satisfy the predictions of the theory of rational choice. This in turn leads us to reconsider our definition of rational behavior.

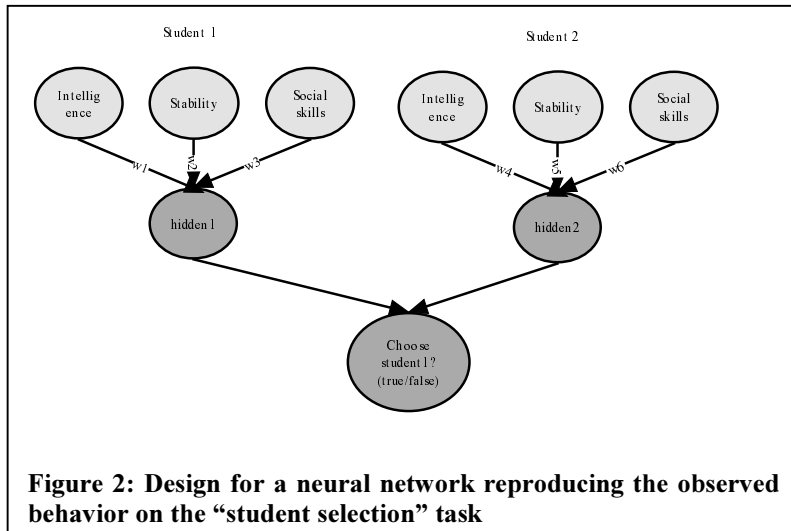
As a thought experiment, imagine the design of two Artificial Neural Networks. The first of these networks is designed to reproduce the *predicted* behavior of volunteers on the "student selection" task described earlier; the second network is designed to reproduce their actual, observed behavior.

The task of reproducing the predicted behavior is a simple one. Any algorithm which computes a single valued score for each student as a function of her intelligence, stability and social skills will automatically demonstrate transitivity. If the results of the tests are coded as integers with maximum value  $k$  even the most complex scoring function can be

represented as a look-up table with  $k^3$  entries. Applying a number of simplifying assumptions we can model our network as the “single layer perceptron” depicted in Figure 1.



**Figure 1: Design for a neural network displaying the predicted behavior on the “student selection” task**



**Figure 2: Design for a neural network reproducing the observed behavior on the “student selection” task**

Modelling the actual behavior is rather more complex. Here the selector has to compare couples (or other groupings) of students. The network will therefore require (at least) twice as much input (intelligence, stability, social skills for student1, intelligence, stability, social skills for student2). For an arbitrary scoring system this requires a look-up table of (at least)  $k^6$  entries. If we try and train a network to reproduce the behavior observed in the laboratory we discover that the problem is not

“linearly separable” and cannot be resolved by a single layer perceptron. The simulation of non-transitive choices requires the kind of multi-layer network shown in Figure 2. Unexpectedly we discover that “imperfect rationality” is harder to model than the “perfect rationality” predicted by the theory of rational choice. Why, one asks, should nature choose a complex solution with poor performance when there exists a simpler solution which would work better? Other work in “New Artificial Intelligence” may help to provide an answer.

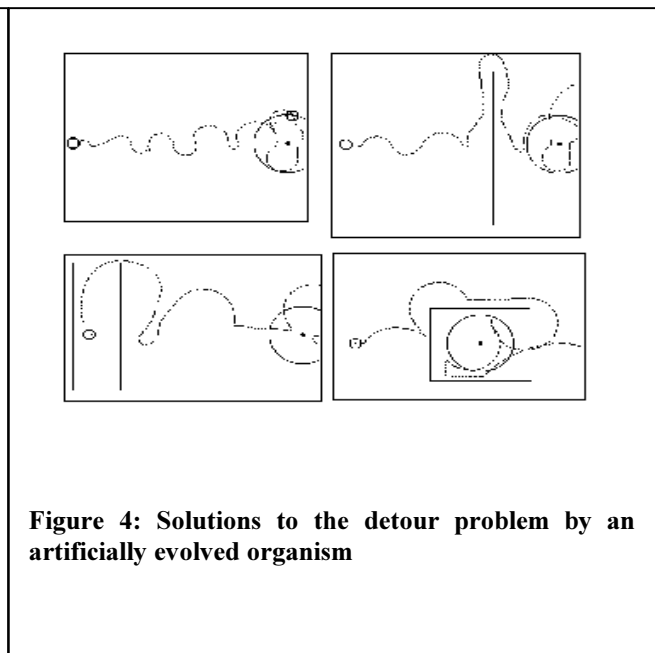

The conceptual roots of Artificial Neural Networks are grounded in an “*information processing*” view of cognition. Neural networks take an input and process this input to produce an output [Rumelhart & McLelland, 1987]. More recent work in the “New Artificial Intelligence” has tended, on the other hand, to take an *ecological stance*, considering not just the neural network but the way in which organisms interact with the environment [Clark, 1997; Parisi, 1999]. In this interaction the only data which the organism can use is the data provided by its sensors. Sensor input depends on the position of the organism and thus on its own actions. Organisms take an *ego-centric view* of their environment; rational behavior consists in making the best possible use of this view.

Ecological models raise a series of new issues for models of cognition and decision-making. One of these is the issue of so-called “Perceptual Aliasing” [Nolfi, 1996]. In an *ecologically grounded* model different environmental situations, requiring different decisions by the organism, may, at a particular moment, generate identical input to the organism’s sensors. A large object, viewed from a distance, may, for example, produce the same image on the retina as a smaller object viewed from closer to. This has implications for our theories of decision-making. In particular it means that *for a physical organism, inhabiting a physical environment there exists no simple mapping from input to output capable of producing biologically optimal behavior. A biologically optimal decision-making mechanism has to be context-sensitive, to react in different ways at*

*different times*. This leads to new questions. How can we (or nature) design an adaptive decision-making mechanism? What are likely to be the characteristics of such a system? Here again “New Artificial Intelligence” can provide us with useful insight.

The decision-making capabilities of natural organisms are the fruit of an evolutionary process. “New Artificial Intelligence” has produced models which simulate this process to evolve “artificial organisms” with specific cognitive abilities [Holland, 1975; Mitchell 1999]. This work has made it possible to understand some of the constraints which apply to evolutionary processes. One of these constraints is that evolution necessarily proceeds “step by step”. In our experiments we have discovered that it is impossible to evolve a complex behavior in a single step. The most effective strategy is to evolve relatively small modules, with relatively simple behaviors, and then to evolve *connections*, between these modules so as to produce more complex, *emergent* behaviors. A second constraint is the need for redundancy. In work currently in progress we show that populations of artificial organisms evolve most effectively when individual organisms in the population are “redundant”, that is when they include many more components than are strictly needed to do the job. Last but not least experimenters have demonstrated that evolution works most effectively when it is combined with *learning*. D. Floreano of EFL Lausanne [Floreano & Urzelai, 1999], for example, has produced models which use evolution to design the *learning rules* for individual synapses allowing the weights of the synapses to change continuously during interactions between the organism and the environment.

Putting these constraints together we can derive a number of consequences for a theory of decision-making. The requirement that evolution should proceed “step by step” implies that it is very hard to evolve big, complex rules in a single step and that it is easier to evolve hierarchies of simple rules, perhaps with complex emergent properties. The requirement for redundancy means that a system governed by a large, functionally inefficient, set of rules, will have a greater ability to adapt to a changing environment than a smaller, more efficient system. Most important of all the requirement for learning signifies that effective decision-making procedures have to be capable of changing over time – perhaps on a very rapid time scale. In short a “biologically rational” decision-making system should make its decisions using large numbers of simple rules based on “local input”; it should, what is more have the ability to change its rules rapidly in response to changes in



context. Systems which meet these criteria are likely to benefit from an important side-effect. It is precisely this kind of system which is likely to prove most robust when faced with environmental change.

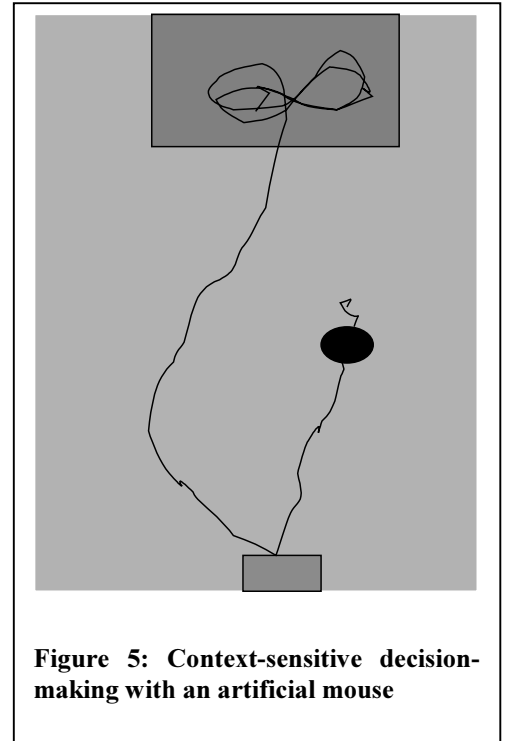
Two examples can help to illustrate the kind of robust decision-making system which evolution can produce. The first of these systems, developed by one of the authors, in collaboration with O. Miglino, [Walker & Miglino, 1999] is a system designed to simulate so-called detour behavior when an animal (see Figure 3) the animal “detours” round an obstacle in order to reach a target. In the experiment the authors “evolved” a population of artificial organisms with the ability to perform the task. The behavior of a typical organism is shown in Figure 4. This behavior can be summarized in three rules:

- •Main: Move forward turning first clockwise (slowly) and then anti-clockwise (more rapidly).
- •Taxis: On visual contact with target turn sharply towards the target. Return to Main
- •Wall following: If left proximity sensors active turn right. If right sensors active turn left. Move forwards until obstacle out of view. Return to Main

At no time in this experiment does the “animal” have a complete view of its environment; at no time therefore is it in a position to evaluate the consequences of its future actions. The information used by the organism comes completely from its sensors and depends not just on the objective characteristics of environment but also on the animal’s position within the environment. This in turn depends on the animal’s previous actions. The “rationality” of the animals behavior derives, in the last analysis, from complex interactions between extremely simple rules, which guarantee that the detour will function robustly in a broad range of different environments.

The second example of evolved rationality is based on the work by D. Floreano, referred to earlier in this paper. In this experiment [Floreano & Urzelai, 1999] Floreano evolved an “artificial mouse” with the ability to solve a simple, context-sensitive decision problem. The “mouse” inhabited the environment shown in Figure 5. In this environment the dark grey area was lit by an electric light controlled by a switch (the small rectangle at the bottom of the diagram). The mouse was rewarded for the amount of time it spent in the illuminated area *when the light was on*. When the light was *off* the optimal behavior for the mouse was to go to the switch and switch it on; when it was on the optimal behavior was to go to the illuminated area and stay there. Floreano successfully evolved mice with the ability to rapidly change their decision rules as the context rewarded different behaviors. The result is the behavior illustrated in the diagram. For these mice the mapping between input and output was a dynamic one. Decision-making was *adaptive*.

Bearing in mind the need to use simple, multiple, highly adaptive rules of decision-making we can now try and model an ideal student selector. Assuming that the selector’s goal is to produce a functional, productive academic community it is easy to see that the application of a static set of decision rules is likely to prove counter-productive. If all students are chosen using the same criteria the result is likely to be a “monoculture” of very similar students all of whom will be very intelligent, or very stable or very sociable, according to the actual decision rules chosen. A rational decision-making procedure should, on the other hand, display adaptivity. If, for example, the current rules are producing an over-concentration of very intelligent, unstable, unsociable students the selection criteria should shift so as to increase the input of stable individuals with better social skills. The skilled decision-maker will, in practice, apply many different rules adapted to particular



**Figure 5: Context-sensitive decision-making with an artificial mouse**

circumstances. In a purely static laboratory setting the application of such rules may not meet abstract criteria for perfect rationality. A decision-maker, taking her decisions in this way, is likely however to be far more adaptable and to produce far better long-term results than a colleague, applying a single mechanical scoring system.

#### **4 CONCLUSIONS: THEORIES OF DECISION-MAKING AND PERCEPTION**

On the basis of what has been said so far it can be argued that the application of traditional criteria for rationality is likely to produce ecologically sub-optimal decision-making. If this is so adaptivity is a fundamental aspect of “rational” decision-making and the “imperfect rationality” observed in the laboratory is a symptom not of *insufficient intelligence* or *imperfect information* but of an artificially static test environment where adaptivity is of unusually low value.

A dynamic theory of rationality and of decision-making has important implications for theories of perception. The goal of “traditional theories”, such as those proposed by Marr [Marr, 1982], is to build systems with the capability to construct a complete, objective representation of external reality. In this framework perception is seen as a computation generating a high level “symbolic” representation of reality from low level sensorial input.

In an ecologically grounded theory of perception on the other hand the goal of perception is to support biologically rational behavior and the generation of rational behavior is seen as an on-going process involving continual adaptation to context via the use of multiple, redundant, very simple computational rules. The kind of theories produced by researchers such as Marr, or traditional theorists in the field of Artificial Neural Networks assume that perception is essentially a static computation. Translated into neuro-biological terms this means that information is transmitted from neuron to neuron by neuro-transmitter release across synapses. Once a network had been “trained” synapses were taken to be static. Perceptive sub-systems of the brain act as a computers calculating output from the motor system as a function of sensorial input. More recent biological research shows however that this view is a serious distortion of reality. Even extremely primitive biological animals display a surprising degree of synaptic plasticity. During the life of the organism the strengths of synapses change continuously on a number of different time scales; the pattern of transmission of information between neurons is subject to continual modification via the action of a large family of neuro-modulators which inhibit or excite the action of particular neural circuits [Katz, 1999]. The biological evidence supports the concept of perception as a dynamic system serving ecological goals. The challenge today is to translate this point of view into concrete theories and models which can provide inspiration both for the biologist and the engineer.

#### **REFERENCES**

Clark, A. (1997), *Being there. putting brain, body and the world Together Again*, MIT Press

Floreano, D. & Urzelai, J. (1999), *Evolution of Neural Controllers with Adaptive Synapses and Compact Genetic Encoding in Lecture Notes in Artificial Intelligence 1674, Advances in Artificial Life*, Springer Verlag

Holland, J.H. (1975), *Adaptation in Natural and Artificial Systems*, MIT Press

Kahneman, D. & Tversky, A (1979), *Prospect Theory: an analysis of decision under risk, Econometrica* 47: 263-291

Katz, P.S., (1999), *Beyond neurotransmission*, Oxford University Press

Marr, D. (1982), *Vision*, W.H. Freeman and Company

Mitchell, M. (1999), *An Introduction to Genetic Algorithms*, MIT Press

Nolfi, S. (1996), Adaptation as a more powerful tool than decomposition and integration in T. Fogarty and G. Venturini (eds.) *Proceedings of the Workshop on Evolutionary Machinery and Machine Learning, 13<sup>th</sup> International Conference on Machine Learning*, University of Bari, Italy, (available at <http://kant.irmkant.rm.cnr.it/nolfipub.htm>)

Parisi, D.,(1999), *Mente: i nuovi modelli della vita artificiale*, Il Mulino

Rumelhart D., McLelland J.L et al., (1987), *Parallel Distributed Processing: explorations in the microstructure of cognition*, MIT Press

Shamir, E. & Tversky, A. (1995), *Decision Making* in Smith E. & Osherson D.N., *An Invitation to Cognitive Science: Thinking*, Volume 3, MIT Press

Tversky, A. (1969), Intransitivity of Preferences, *Psychological Review*, 76: 31-48

Tversky, A., & Kahneman, D. (1981), The framing of decisions and the psychology of choice. *Science*, 211, 453-458

Walker, R. & Miglino, O., (1999), Replicating Experiments in “Detour Behavior” with Artificially Evolved Robots: an A-Life Approach to Comparative Psychology