Many-Agent Simulation and Artificial Life
E. Hillebrand and J. Stender (Eds.)
IOS Press, 1994

Artificial Intelligence and Economic Theory

Nicolaas J. VRIEND*
Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA, <vriend@santafe.edu>

Abstract. Recently, economists have shown a rapidly growing attention for the field of artificial intelligence (AI). This contribution does *not* discuss the *technology* of AI, or its *applications* to econometrics, business, finance or management. Instead, we explain the significance of AI for *economic theory*; in particular for the theory of decentralized economies.

Are you after truth? Yeah. But I don't know what we mean by truth in our business. I don't see economics as pushing that deeply in some respects. We're programming robot imitations of people, and there are real limits on what you can get out of that. (Lucas in [26], p. 49)

1. Introduction

Recently, economists have shown a rapidly growing attention for the field of artificial intelligence (AI). This contribution does *not* discuss the *technology* of AI, or its *applications* to econometrics, business, finance or management. Instead, we explain the significance of AI for *economic theory*; in particular for the theory of decentralized economies. In section 2 we expose the essence of economic theory, showing how Lucas' assertion that doing economics implies *"programming robot imitations of people"* (see motto) was meant as a metaphor. Section 3 is a digression on AI, while section 4 considers the employment of AI in economic theory, arguing that the current availability of AI techniques makes it worthwhile to take Lucas' observation literally.

* I wish to thank Brian Arthur, Raja Das, Pierre Dehez, Alan Kirman, and Martin Shubik for critical comments and helpful discussions on these issues. All errors and responsibilities are mine. The Santa Fe Institute's stimulating environment, and a grant from the Niels Stensen Foundation are gratefully acknowledged.

2. Economic Theory

2.1 Fundamentals

It is widely accepted that the science of economics started with Adam Smith. The main accomplishment of Smith was to put into the center of economics the systematic analysis of the behavior of individual agents pursuing their self-interest under conditions of competition. The most eloquent quotation in this respect is presumably: *"It is not from the benevolence of the butcher, the brewer, or the baker, that we expect our dinner, but from their regard to their own interest"* ([40], p. 26/27). Since then, this axiom concerning the behavior of individual agents has, as a matter of course, become a fundamental part of economic discourses.[1]

A century later Edgeworth [15] considered it useful to articulate this very explicitly and precisely: *"The first principle of Economics is that every agent is actuated only by self-interest"* (p. 16). To appreciate this assertion of Edgeworth fully, it may be necessary to examine this compound statement carefully. The second part asserts something about individual agents which echoes Smith. The ultimate motive for any action must be found in the agent's desire, agents are acting only out of self-interest. This presupposes that it is evident what is meant by the term *self-interest*. Edgeworth [15], more than a century ago, used the word *"pleasures"*, defined as *" 'preferable feeling' in general"* (p. 56). In the language of present-day economic discourses, what is self-interest is a matter of *preferences*. Next, let us consider the first part of Edgeworth's assertion. He claims that this is the first principle, the starting-point, of economics. In other words, the statement about individual agents driven exclusively by self-interest is a defining statement concerning the *homo oeconomicus*. The *homo oeconomicus* is an agent with given preferences.

Given these preferences, the *homo oeconomicus*, pursuing his self-interest, seeks to do the best he *can*. That is, it is important to pay explicit attention to the *homo oeconomicus*' opportunities and his perception of these opportunities. Perceived opportunities are perceived possible actions plus perceived consequences. These perceptions themselves depend on economic behavior. First, as information is a valuable asset, the information that an individual agent has, in particular his perception of opportunities, is the result of economic behavior (see [41]). Secondly, also the development of cognitive skills is a result of economic behavior (see [9]). Thus, opportunities are defined such that *all* perceived costs and benefits are taken

[1] Whether this was also exactly as Smith himself intended to put these matters is an interesting, but *different*, question (see, e.g., [23]).

into account; in particular information, decision making and transaction costs. Opportunities are not necessarily only transaction opportunities. Agents may also have possibilities to search, to talk with a friend, to go to school or to the beach, to do nothing, etc. This is most clearly stated by Becker [10]: *"When an apparently profitable opportunity ... is not exploited, the economic approach does not take refuge in assertions about irrationality .... Rather it postulates the existence of costs, monetary or psychic, of taking advantage of these opportunities that eliminate their profitability - costs that may not be easily "seen" by outside observers"* (p. 7).

Economic behavior simply means that an individual agent chooses (one of) the most advantageous options, given his preferences, in his perceived opportunity set. Hence, given the *homo oeconomicus*' perceived opportunities and preferences, his actions can be derived rather *mechanically*. It is this what Lucas meant when asserting that doing economics is like *"programming robot imitations of people"* (see motto).

2.2 *Modeling the homo oeconomicus*

Having established that the *homo oeconomicus*' actions depend upon his preferences and perceived opportunities, the central concern is how to model this. One way to deal with preferences in economic theory would be to ask advice about their properties from, for example, psychologists. However, one could wonder why economists would bother much to make specific assumptions concerning individual preferences, even if one would agree that these preferences drive the individual's actions. Until recently, the idea was the following. By making assumptions about individual preferences one wanted to derive certain characteristics of aggregate behavior. By now we know that it is theoretically impossible to get needed characteristics of aggregate demand functions (needed in order to prove stability of the tâtonnement process) by imposing more and more restrictions upon individual characteristics (see [25] for a survey). In other words, in the aggregate, the assumptions of individual preferences have in general no implications (see also [2]). Therefore, approaches which rely less upon specific assumptions concerning individual preferences may be more promising. Stigler and Becker [42] argue that preferences should not only be taken for given in economics, but can also be considered roughly the same for everybody. Differences in actions are then completely ascribed to differences in perceived opportunities. Still further goes Becker's [8] exercise, which focusses exclusively upon the perceived opportunity set.

Allowing for virtually every imaginable type of individual behavior,[2] he analyzes the relations between opportunity sets of individual agents and market outcomes.

This points to the second important problem concerning economic models: the modeling of the agents' perceived opportunities, without turning economics into a psychology of perception. Basically, the problem is that economists are definitely not in a position to contribute to an explanation of how a set of given physical stimuli, including both the agent's objective environment and his own brain status and activity, leads to a set of perceived opportunities. When, in the economic process, perceived opportunities evolve over time, these changes will not only be due to a change in the perception of the underlying circumstances, i.e., learning, but also to a change in these circumstances themselves, as a result of the interactions between the agents. And, in general, these learning processes and the other dynamic economic forces may interact with each other. This points to the following way to abstract from psychological matters concerning the perception of opportunities.

Assume that the perception of opportunities is an endogenous process. That is, the set of perceived opportunities depends strictly upon the preceding sequence of actions and outcomes. While the agents' actions depend on their perceived opportunities, these opportunities and their perceptions depend on the agents' own market experience as the result of previous actions. Thus, in a formal model, actions will be a function of perceived opportunities, *and* perceived opportunities a function of earlier actions. As a result one gets a sequence analysis of actions as function of previous actions and outcomes, while perceptions or expectations do not appear explicitly but only *"between the lines"* ([18], p. viii).

The crucial issue, then, is the specification of such functions, mapping the agents' past actions and outcomes into current actions. Clearly, to tie down the set of functions a priori in an ad hoc way, assuming simple fixed rules-of-thumb, would not be very interesting. In the next section we will show how the current availability of artificial intelligence techniques may be useful here.

3. Artificial Intelligence

In this section we discuss three approaches to machine learning that may be relevant to economic theory. We will not argue that rational economic agents do use such AI techniques; the *'as if'* argument will do. The approaches examined are: Genetic Algorithms, Classifier Systems, and Artificial Neural Networks. Both CSs and GAs

---

[2] Becker [8] calls it *'irrational'* behavior, which he defines as every kind of behavior *not* equal to choosing the most preferred option in the perceived opportunity set.

have for a large part been developed in the *'school'* of John Holland at Ann Arbor, Michigan (see, e.g., [20], [21] and [22]). Useful introductory surveys can be found in [14] or the special issue on genetic algorithms of Machine Learning [30]. [17] is an excellent elementary but comprehensive textbook. For an introduction to ANNs see, e.g., [28] and the references therein. As AI in economics is just a tool, used in order to model individual agents, the presentations of these approaches in this paper serve only pedagogical goals, and are not intended as an exhaustive historical survey or critical discussion.

3.1 Artificial Neural Networks

Artificial Neural Networks (ANNs) are often considered as black boxes. For our purpose, such an approximate view will suffice. ANNs map a set of input features to a set of output features. In order to be able to achieve such a task, an ANN needs some learning. To start with, one needs a training set consisting of a number of input patterns x plus attached to each observation the corresponding *'true'* or *'correct'* value of some output variable y. The input patterns are presented to the ANN, and for each input pattern the ANN's actual output $\hat{y}$ is compared with the correct or *'target'* output y.[3] When the whole batch of input patterns is processed, the internal parameters of the ANN are adjusted on the basis of the errors, which are the differences between the outputs determined by the ANN $\hat{y}$ and the target outputs y.[4] This process is repeated, using the same set of input patterns, until the error is smaller than some given limit.

The most interesting feature of ANNs is that they use some sort of general flexible functional form, without any pretensions about the internal representations of reality, data generating processes, or causal chains, in order to yield an inherently misspecified approximation of an unknown function.[5] Conceptually, *'training'* or *'learning'* with an ANN seems *equivalent* to running an Ordinary Least Squares regression. Given a number of observations concerning some explanatory variables x (input) and the corresponding actual values of a dependent variable y (target output), one calculates parameter values to determine the estimated dependent variable $\hat{y}$ (output) such that some error term, measuring the difference between y

[3] In an economic model, the ANN's input would be the agent's market experience, the ANN's actual output $\hat{y}$ would be its action chosen for the next period, while the *'correct'* or target output would be that action that would maximize payoff.

[4] The most commonly applied method to adjust the parameters is backpropagation (see [47]).

[5] Lippmann [28] refers to a theorem proven by Kolmogorov and described in [29] which effectively states that a three layer ANN with $n(2n+1)$ nodes using continuously increasing nonlinearities can compute *any* continuous function of n variables (see also [48]).

and $\hat{y}$, is minimized. The adapted parameter set or estimated coefficients can then be used to make predictions. Hence, to an econometrician ANNs are a useful new technique to cope with the problem of misspecification.

There are, however, some conceptual problems with the ANN learning method sketched above. The main problem is that the method relies completely upon some external supervisor. In essence, by correcting parameters on the basis of some error function representing a measure of the distance between the *'target'* output y and the ANNs actual output $\hat{y}$, the external supervisor teaches the ANN to *reproduce* the target output for each input pattern in a training sequence. In other words, *'learning'* by such ANNs means generalizing, summarizing and memorizing a *given* input-output mapping.[6] In general, however, and in particular in a decentralized economy, there is no external supervisor to teach the ANN which is the *'correct'* (i.e., the best possible) output, or how much it differed from such a target, not even afterwards. Often, there is only a notion of what the ANN should accomplish plus a success measure of its performance.[7] That is, the ANN has to learn through *'reinforcement'* (see, e.g., [31]). Sometimes the sketched process of error correction in supervised ANNs is also called reinforcement learning. As Barto et al. [7] point out, that is misleading. Error correction mechanisms are not based on a relative assessment of *consequences* of the ANN's output, but *only* on knowledge of the supervisor of both the correct and actual output. This does *not* involve feedback that passes through the ANN's environment. Before we sketch how reinforcement learning can be implemented in ANNs, we give an example to illustrate this important problem further.

According to Zermelo's Theorem: *"In chess either white can force a win, or black can force a win, or both sides can force at least a draw"* ([6], p. 1). Hence, the learning task concerning chess is clear cut: Discover which of these three options apply, and determine the corresponding moves to play. Although chess is an extremely simple game when compared with real life, and although it is even closed

---

[6] One could even question whether supervised ANNs belong to the domain of AI. The most commonly used implicit definition of intelligence applied to AI follows from the *'Turing test'*: If a computer behaves in a conversation in a way as to be confused with a human being, then it should be defined intelligent. This definition leaves no role whatsoever for the role of learning. Poggio [34] reports on a recent experiment in which some very simple computer programs turned out to confuse people (see also [38]), and argues that a system should be considered intelligent when it is able to learn unsupervised. There do exist ANNs that learn without supervision. Usually, these produce classifications simply clustering input data. For example, assuming all handwritten *b*'s look more like each other than like *c*'s or *d*'s etc., they put all *b*'s together in one class, all *c*'s in another class, etc. Afterwards, one just has to label the right class *'B'*, *'C'*, etc.

[7] For example, the ANN has to generate profits or utility, mapping an observed state (input) to actions (output), where the measure of success is simply the amount of profits or utility.

in the sense that the number of possible moves is finite and countable, the number of possible moves and positions exceeds all existing computing power. Nevertheless, a learning ANN would need a measure for the distance between its own evaluation of positions and the correct or *'target'* evaluation of positions in order to adjust its parameters. The makers of Deep Thought, one of the best computer chess players, have resolved this problem in the following way. In some cases the correct evaluations can be found by performing depth first searches. In other cases, they use a batch of 900 master games, and simply define the moves played by these first-rate human players as the optimal or correct moves.[8] Now, by summarizing and memorizing the knowledged expressed by these grandmasters, Deep Thought has caught up with the best human players, and will, perhaps, be able to overtake even Kasparov, actually the best *'supervisor'* available, but this falls well short of learning the game of chess as stated above.[9]

This example shows that the problem with supervised ANNs is *not* that they use information supplied by another agent. A priori there is no reason to distinguish between knowledge based on information about what other agents have done in a certain situation, and knowledge based on own prior experience in such circumstances; and often the former source of knowledge will be much less costly (see [49]). The problem with supervised ANNs is that the knowledge of some other agent is proclaimed *'true'* or *'correct'*. Hence, such a ANN does not much more than trying to imitate a supervisor that is *presumed* to be perfect. In the case of Deep Thought this presumption is clearly inaccurate.

One way to solve the problem of reinforcement learning is using two ANNs (see, e.g., the seminal [7]). The basic ANN gets its input x (the observed state) and produces an action $\hat{y}$ as its output. Some unknown system, e.g., *'the economy'*, then determines a final outcome V. In order to adjust the parameters of the basic ANN such that the *unknown* optimal or *'target'* action y is approximated, i.e., such that the final outcome V as a measure of success will be maximized, one needs information about this unknown system. This information can be constructed as follows. A second ANN learns to mirror the unknown system, mapping the observed inputs x directly to outputs $\hat{V}$ that are a prediction of the actual final outcome V of the system. The target output of this second ANN is the actual V as realized by the unknown system. Learning of this ANN takes place through an error correction

[8] *"[A]ny position reached after a grandmaster's move is, after all, likely to be better than all of the others that would have been reached via alternative moves"* ([24], p. 48/49). Note that this is exactly Friedman's [16] selection argument in his side-remarks about optimizing billiard players.

[9] Although this judgement may seem rather cynical, the makers of Deep Thought themselves are well aware of these limitations: *"Deep Thought ... remembers everything but learns nothing ..."* ([24], p. 50).

mechanism aimed at minimizing the difference between V and $\hat{V}$. Remember that this second ANN does not need to understand the underlying mechanisms of the economic processes which determine the actual outcome V. This second ANN, then, supplies the necessary reinforcement signals to guide the adjustment of the parameters of the basic ANN.

Another, more recent but closely related, approach to the problem of reinforcement learning is the technique of Q-learning [46], in which an ANN is used not only to evaluate the consequences of its actions, both in terms of immediate rewards and its estimate of the value of the state to which it is taken, but also to decide upon the actions (see also [31]).

This discussion of ANNs should have made clear how essential the difference between reinforcement learning and supervised learning is. Understanding this issue helps to see why Classifier Systems and Genetic Algorithms may be useful tools to overcome this obstacle.

3.2 Genetic Algorithms

A Genetic Algorithm (GA) consists of a set of actions, with to each action attached a measure of its strength. This strength depends upon the outcome or payoff that would be generated by the action. Each action is decoded into a string. Through the application of some genetic operators new actions are created, that replace weak existing ones. GAs are search procedures based on the mechanics of natural selection and natural genetics. The set of actions is analogous to a population of individual creatures, each represented by a chromosome with a certain biological fitness. The basic GA operators are *reproduction*, *crossover* and *mutation*. Reproduction copies individual strings from the old to a new set according to their strengths, such that actions leading to better outcomes are more likely to be reproduced. Crossover creates a random combination of two actions of the old set into the new one, again taking account of their strengths. This makes that new regions of the action space are searched through. Mutation is mainly intended as a *'prickle'* every now and then to avoid the set to lock in in a sub-space of the action space. It randomly changes codes of a string, with a low probability.

The key feature of GAs is their ability to exploit accumulating information about an initially unknown search space, in order to bias subsequent search efforts into promising regions, and this although each action in the set refers to only one point in the search space. An explanation of why GAs work is condensed in the so-called

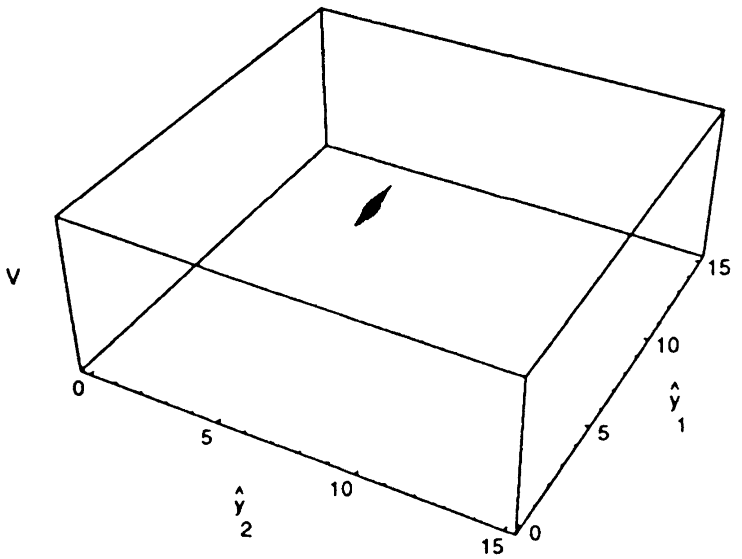

Figure 1 a single point in an unknown landscape

'Schema Theorem'.[10] When one uses the binary alphabet to decode the actions, then 10110*** would be an example of a *'schema'*, where * is a so-called *'wild card'* symbol, i.e., * may represent a 1 as well as a 0. The following example shows the power of these schemata. Suppose an individual agent has two decision variables ($\hat{y}_1$ and $\hat{y}_2$) and an unknown payoff function (V), then the search space may be represented by the metaphor of an unknown landscape. Each action as such refers to only one point in this unknown landscape. As figure 1 shows, this does not contain much information as to where to find the most attractive regions.

Using the binary alphabet and constructing the string by alternating the bits for $\hat{y}_1$ and the bits for $\hat{y}_2$, the portrayed sample action $(\hat{y}_1,\hat{y}_2) = (12,4)$ would be represented by the string 10110000.[11] Hence, 10110*** would be one of the schemata present in this action. This schema contains much more information about the landscape, as figure 2 shows, where the shaded areas are those regions in which all possible combinations of $\hat{y}_1$ and $\hat{y}_2$ are processed implicitly by the genetic operators.

Reproduction, crossover and mutation select strings and then operate on the coded information represented in these strings. Hence, the more the information referring to a single point in the search space is fragmented into small pieces, the more schemata are processed implicitly, and the more information is used by these genetic operators. This leads to the requirement of using the smallest possible decoding alphabet. Not all schemata are processed equally usefully, and many of them will be disrupted by the genetic operators; in particular by the crossover

[10] Also called *'Fundamental Theorem of Genetic Algorithms'* (see, e.g., [17] or [44]).

[11] The string for $\hat{y}_1$ would be 1100, and for $\hat{y}_2$ 0100.

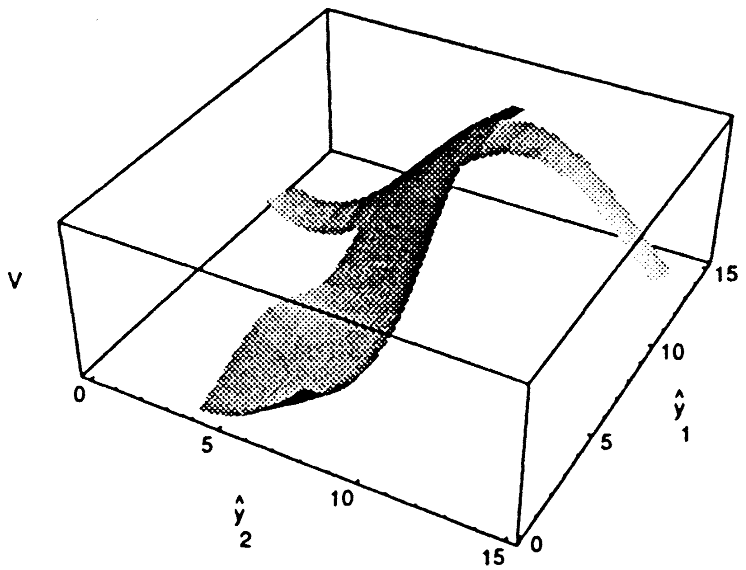

Figure 2 a schema in an unknown landscape

operator. The *'Schema Theorem'* says that short, low-order, high performance schemata will have an increasing presence in subsequent generations of the set of actions, where the order of a schema is the number of positions defined in the string, and the length is the distance from the first to last defined position. Although this *'implicit parallelism'* is also sometimes called *'randomized parallel search'*, this does not imply directionless search, as the search is guided towards regions of the action space with likely improvement of the outcomes.

GAs are especially appropriate when, for one reason or another, analytical tools are inadequate, and when point-for-point search is unfeasible because of the enormous amount of possibilities to process, which may be aggravated by the occurrence of non-stationarity. But the most attractive feature of GAs is that they do not need a supervisor. That is, no knowledge about the *'correct'* or *'target'* action, or a measure of the distance between the coded actions and the *'correct'* action, is needed in order to adjust the set of coded actions of the GA. The *only* information needed are the outcomes that would be generated by each action. In this sense GAs exploit the local character of information, and no further knowledge about the underlying outcome generating mechanisms is needed, like e.g., the derivatives of certain functions.

Although a GA does not need information concerning the *'correct'* action, a drawback of GAs is that they still do need, for every coded action present in the set, the information concerning the outcome that would be generated by *that* action. When there is no supervisor, typically not even such information will be available. Note that this information requirement is considerably less than in the case of a supervised ANN. We will now examine Classifier Systems, and show that those can be used to supply the necessary information by implicitly constructing a prediction of the outcomes for all actions in the set.

3.3 Classifier Systems

A Classifier System (CS) consists of a set of decision rules of the *'if ... then ...'* form. To each of these rules is attached a measure of its strength. Actions are chosen by considering the conditional *'if ...'* part of each rule, and then selecting one or more among the remaining rules, taking into account their strengths. The choice of the rules that will be activated is usually determined by means of some stochastic function of the rules' strengths.

The fundamental virtue of CSs is that it aims at offering a solution to the reinforcement learning or *'credit assignment'* problem. A complex of external payments and mutual transfers of fractions of strengths can be implemented, such that eventually each rule's strength forms implicitly a prediction of the payoff it will generate when activated. The basic source from which these transfers of strengths are made is the external payoff generated by an acting rule. The strengths of rules having generated good outcomes are credited, while rules having generated bad outcomes are debited. Thus the outcomes encountered *'induce'* successive actions. Note that one can distinguish two levels of endogenity in a CS. First, the set of 'if ... then ...' rules forms explicit links between states and actions, i.e., between the outcomes of previous actions and subsequent actions. Secondly, the strengths of these relations between states and actions develop endogenously, i.e., the relative strengths of the rules in the set are determined by the rules actually executed and by the outcomes they have actually generated.[12] Two factors make that the direct reward from the CS's environment to the acting rule does not necessarily reinforce the right rules. First, the state in which the CS happens to be may depend, among other things, upon previous decisions. This is important, as only those rules of which the conditional *'if ...'* part was satisfied could participate in the decision of the current action. Hence, when the current decision turns out to give high payoffs, it may be the rules applied in the past which gave that rule a chance to bid. An example is the game of chess, where the final move, the one that actually receives the payoff from the environment, can be made only thanks to numerous preceding moves. Secondly, more in general, it may be that not all payoffs are generated immediately, due to the presence of lags or dynamics, implying that the current outcomes are not only determined by the current action, but also partly by some actions chosen previously. This credit assignment problem is dealt with by the so-called *'Bucket Brigade Algorithm'*. In this algorithm each rule winning the right to be active makes a payment to the rule that was active immediately before it. When

---

[12] This endogenity is the main difference between CSs and Expert Systems, where these links are determined a priori by the expertise of the creator of the system.

the CS repeatedly goes through similar situations, this simple passing-on of credit makes that the external payoff may be distributed appropriately over complicated sequences of acting rules leading to payoff from the environment.[13]

Note that Classifier Systems and GAs are complementary, and they can very well be applied as a combination.[14] While CSs are used to govern the reinforcement learning process, determining the strengths of the actions and determining which action will actually be executed, the GAs can be used to generate new sets of actions. The frequency at which the latter is used is determined by the GA rate. Note that a too high GA rate would make that the CS does not get enough time to predict the value of the newly created strings, while a too low GA rate would lead to lack of exploration of new regions.

4. The Significance of AI for Economic Theory

We have seen in section 2 that the fundamental characteristic of the *homo oeconomicus* is that he just chooses the most preferred option in his perceived opportunity set. We have also argued how the need for abstraction from psychological issues concerning the perception of opportunities, led to the idea of a sequence analysis of actions as functions of previous actions and outcomes. The property that makes the CS/GA approach so fruitful for economic theory is, that the relations between actions and previous actions and outcomes can be kept completely *flexible*. This implies that one is in a position to analyze how far *'the market'* provides sufficient structure to tie down the set of perceived opportunities, i.e., to constrain the behavior of the individual agents (cf., [8]). This is what one could call, following Blume and Easley [13], a *'positive theory of action'*. Hence, nothing seems more obvious than taking Lucas' assertion that doing economics implies *"programming robot imitations of people"* (see motto) literally. Therefore, one could run a many agent simulation of a decentralized economy, in which each individual *homo oeconomicus* is programmed separately applying a CS/GA, seeking to do the best he can in his unknown payoff landscape. As the individual agents interact with each other, these landscapes for the individual agents may co-evolve.[15]

---

[13] For an analysis of the similarities between the *'Bucket Brigade Algorithm'* and the method of backpropagation used in ANNs, and between CSs and ANNs in general, see [11].

[14] Often GAs are presented as an add-on to CSs, or the other way round. However, although CSs and GAs are closely related to each other, it seems useful to distinguish them conceptually very clearly.

[15] See, e.g., [32], [1], [36], [5], or [45].

Note that the agents modeled with a CS/GA are not *'myopic'*. In a CS/GA the whole history of the agents' experience counts, and they are competent enough, to give up direct profits/utility, in order to gather information to generate more payoff later on. Moreover, also rules that do not directly generate payoff are reinforced according to their merits. This makes that agents may *'recognize'* valuable *sequences* of actions.

It would also be confusing to depict the behavior of the individual agents modeled by a CS/GA as *'adaptive'*, and it might be evidence of an important misconception of the issues at stake. Typically, *'adaptive'* behavior is thought to mean something as *'too passively walking behind the facts'*. Such a description would be fully inappropriate for the agents modeled by a CS/GA. These agents are active searchers for the most advantageous opportunities. They experiment to improve their perceptions of these opportunities, continuously exploring the most promising regions of their action domain. The crucial point is that what the agents perceive to be promising is a function of the exogenously given information at the start of the process, and all the experiences during the process. What is excluded are ad hoc *exogenous* changes of perceptions during the process, because those would sweep away every hope to find constraints imposed by the market process upon the individual agents' possibly perceived opportunities.

It should also be stressed that the CS/GAs are not models of agents using only simple decision rules. Although each rule for itself in a CS/GA is a simple rule, it is the *set* of rules that forms the link between actions and previous actions and outcomes, and it is not the individual rules that matter. Moreover, this set of rules may change, applying the genetic operators. As is well-known, such a representation of knowledge is not restrictive in any sense, and any program that can be written in a standard programming language can be implemented in a CS.[16] Hence, a CS/GA may be thought to model the most complex and sophisticated human decision procedures, as well as the most simple. In other words, *any* decision can be modeled *'as if'* made by a CS/GA.

Two possible criticisms of many agent simulations using CS/GAs might be that the behavior of the agents is *ad hoc* and the way the agents are modeled is *arbitrary*. Both would be correct observations, but, as we will argue here, only in the following very specific sense.

A general characteristic of agents living in the complexity of a *'large world'* is that they do not have a true, well-specified model to work with. That is, the agents' problem situation is ill-defined (see [3] and [4]). Hence, instead of basing their

---

[16] That is, these systems are *'computationally complete'* (see [33]).

actions on deductive reasoning from universal truths, they are forced to inductive reasoning. Inductive reasoning proceeds from the actual situation faced by an agent. In this sense, such agents' behavior is adaptive or reactive. Sometimes this is also known as the *'cross that bridge when you come to it'* principle (see [37]), because "*... in a large world ... there are some bridges that you cannot cross before you come to them*" ([12], p. 1). Hence, it is only in a very *literal* sense that inductive behavior might be called *'ad hoc'*. Note that it does not imply in any sense an *'anything goes'*, i.e., an abandoning of logical principles or rationality. It would seem to come close to rationality in the sense of *'situational logic'* (see, e.g., [19] or [35]).

Modeling this inductive behavior of the individual agents with CS/GAs is certainly arbitrary, but any approach would be arbitrary *to some extent*. Remember that, in general, in a decentralized economy the agents cannot perceive what the objectively optimal actions would be. We have argued that economists do not have the tools to construct explicit mental models for the agents' perceptions, and that hence, we could follow the approach of mapping actions and outcomes directly to new actions, leaving the mental processes implicit. This mapping, in order to determine the agents' new actions, is not fixed a priori, but kept flexible. Competing hypotheses are tested and their perceived usefulness is updated in parallel. Reinforcement of hypotheses takes place on the basis of payoffs experienced in the market. New hypotheses are formed from building blocks of rules that had turned out to be useful. Bad hypotheses are easily discarded as experience accumulates. Thus, reinforcement, through actual payoffs experienced in the market plays, the pivotal role in a CS/GA. This means that as far as these algorithms are arbitrary, it is the market that acts as the arbitrator! For an economist that must be more than reasonable.

Although CS/GAs are not the only possible algorithms in this context, it seems that alternative algorithms have to meet at least the following three requirements. Firstly, they should be equally flexible as to the possible mappings from the agents' previous actions and outcomes to current actions. Secondly, the market should play an equally essential role in directing the agents' actions. Thirdly, the dynamics of learning and the dynamics of the economic forces as such should be modeled at the appropriate two, conceptually distinct, levels.[17]

In order to answer the question whether the market provides sufficient structure, one has to look for the emergence of regularities in the actions and outcomes during the process of creating and trading away of opportunities by economic agents. Interesting are those regularities that cannot be deduced directly from the built-in

---

[17] Cf., many *'evolutionary'* models in economics in which some form of *'replicator dynamics'* is applied, modeling these two types of processes at the same, population level.

properties of the individual agents or some other microeconomic aspect of the model; at least not by any argument which is substantially shorter than producing that regularity by running the simulation itself (see [27]). The emergence of such regularities is usually related to the metaphor of the *'Invisible Hand'*. While the individual agents take care only about their own self-interest, it is the *'Invisible Hand'* that is thought to perform a regulating function, bringing about coordination of economic activities.

The final objective of such type of analysis is not to become wise with respect to artificial worlds, but to understand what is going on in *real* decentralized economies. Therefore, a serious question to examine would be, whether it is possible to *'recover'*[18] regularities known from reality in, necessarily simple, simulated models, and to analyze how these regularities depend upon parameter choices or modeled mechanisms. Simulations of artificial economies fulfill here the same role as any formal, mathematical model that abstracts from some aspects of reality. They may suggest ways how one might understand what is going on in a decentralized economy.

References

[1] Andreoni, J., & Miller, J.H. (1991). *Auctions with Adaptive Artificially Intelligent Agents* (Working Paper 91-01-004) Santa Fe Institute.

[2] Arrow, K.J. (1986). Rationality of Self and Others in an Economic System. *Journal of Business, 59, No. 4, Pt. 2*, S385-S399.

[3] Arthur, W.B. (1992). *On Learning and Adaptation in the Economy* (Working Paper 92-07-038) Santa Fe Institute.

[4] Arthur, W.B. (1994). Inductive Reasoning and Bounded Rationality. *American Economic Review, 84, Papers and Proceedings.*

[5] Arthur, W.B., Holland, J., LeBaron, B., Palmer, R., & Taylor, P. (1994). *An Artificial Stock Market* (mimeo).

[6] Aumann, R.J. (1976). *Lectures on Game Theory* (Underground Classics in Economics) Stanford University.

[7] Barto, A.G., Sutton, R.S., & Anderson, C.W. (1983). Neuronlike Adaptive Elements That Can Solve Difficult Learning Control Problems. *IEEE Transactions on Systems, Man, and Cybernetics, 13, No. 5*, 834-846.

[8] Becker, G.S. (1962). Irrational Behavior and Economic Theory. *Journal of Political Economy, 70*, 1-13.

[9] Becker, G.S. (1964). *Human Capital: A Theoretical and Empirical Analysis.* New York: Columbia University Press (NBER).

[10] Becker, G.S. (1976). The Economic Approach to Human Behavior. In G.S. Becker, *The Economic Approach to Human Behavior* (pp. 3-14). Chicago: Chicago University Press.

[11] Belew, R.K., & Forrest, S. (1988). Learning and Programming in Classifier Systems. *Machine Learning, 3, Nos. 2/3*, 193-223.

[18] Cf., the notion of *'recoverability'* in [43].

[12] Binmore, K.G. (1991). *DeBayesing Game Theory* (mimeo) Lecture for the International Conference on Game Theory, Florence.

[13] Blume, L.E., & Easley, D. (1991). *What Has the Learning Literature Taught Us?* Lecture given at the Summer Workshop on Learning and Rationality in Economics, European University Institute, Florence.

[14] Booker, L.B., Goldberg, D.E., & Holland, J.H. (1989). Classifier Systems and Genetic Algorithms. *Artificial Intelligence, 40*, 235-282.

[15] Edgeworth, F.Y. (1881). *Mathematical Psychics*. London: Routledge.

[16] Friedman, M. (1953). The Methodology of Positive Economics. In M. Friedman, *Essays in Positive Economics*. Chicago: University of Chicago Press.

[17] Goldberg, D.E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley.

[18] Hart, A.G. (1951). *Anticipations, Uncertainty and Dynamic Planning* (2nd ed.). New York: Kelley.

[19] Hayek, F.A. (1973). *Law, Legislation, and Liberty* (Vol. 1, Rules and Order). Chicago: University of Chicago Press.

[20] Holland, J.H. (1986). Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems. In R.S. Michalski, J.G. Carbonell & T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach (Vol. 2)* (pp. 593-623). Los Altos, CA: Morgan Kaufmann.

[21] Holland, J.H. (1992). *Adaptation in Natural and Artificial Systems. An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (2nd ed.). Cambridge, MA: MIT Press.

[22] Holland, J.H., Holyoak, K.J., Nisbett, R.E., & Thagard, P.R. (1986). *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.

[23] Holmes, S. (1990). The Secret History of Self-Interest. In J.J. Mansbridge (Ed.), *Beyond Self-Interest* (pp. 267-286). Chicago: University of Chicago Press.

[24] Hsu, F., Anantharaman, T., Campbell, M., & Nowatzyk, A. (1990). A Grandmaster Chess Machine. *Scientific American, 263, No. 4*, 44-50.

[25] Kirman, A.P. (1992). Whom or What Does the Representative Individual Represent? *Journal of Economic Perspectives, 6, No. 2*, 117-136.

[26] Klamer, A. (1984). *The New Classical Macroeconomics. Conversations with New Classical Economists and their Opponents*. Brighton: Wheatsheaf.

[27] Lane, D.A. (1992). *Artificial Worlds and Economics* (Working Paper 92-09-048) Santa Fe Institute.

[28] Lippmann, R.P. (1987). An Introduction to Computing with Neural Nets. *IEEE ASSP Magazine, April*, 4-22.

[29] Lorentz, G.G. (1976). The 13th Problem of Hilbert. In F.E. Browder (Ed.), *Mathematical Developments Arising from Hilbert Problems* (pp. 419-430). Providence, RI: American Mathematical Society.

[30] *Machine Learning* (1988). *3, Nos. 2/3, Special Issue on Genetic Algorithms.*

[31] *Machine Learning* (1992). *8, Nos. 3/4, Special Issue on Reinforcement Learning.*

[32] Marimon, R., McGrattan, E., & Sargent, T.J. (1990). Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents. *Journal of Economic Dynamics and Control, 14*, 329-373.

[33] Minsky, M.L. (1967). *Computation: Finite and Infinite Machines*. Englewood Cliffs, NJ: Prentice-Hall.

[34] Poggio, T. (1992). Caro signor Turing urge un altro test. E una macchina che impari da sola. *La Repubblica, 30 settembre, supplemento*, 3.

[35] Popper, K.R. (1957). *The Poverty of Historicism*. London: Routledge & Kegan Paul.

[36] Rust, J., Palmer, R., & Miller, J.H. (1992). *Behavior of Trading Automata in a Computerized Double Auction Market* (Working Paper 92-02-008) Santa Fe Institute.

[37] Savage, L.J. (1954). *The Foundations of Statistics*. New York: Wiley.

[38] Shieber, S.M. (1994). Lessons from a Restricted Turing Test. *Communications of the ACM, 37, Nr. 6,* 70-78.

[39] Simon, H.A. (1986). Rationality in Psychology and Economics. *Journal of Business, 59, No. 2, Pt. 2,* S209-S224.

[40] Smith, A. (1976). *An Inquiry into the Nature and Causes of the Wealth of Nations. (Glasgow edition).* Oxford: Clarendon.

[41] Stigler, G.J. (1961). The Economics of Information. *Journal of Political Economy, 69,* 213-225.

[42] Stigler, G.J., & Becker, G.S. (1977). De Gustibus Non Est Disputandum. *American Economic Review, 67,* 76-90.

[43] Varian, H.R. (1984). *Microeconomic Analysis* (2nd ed.). New York: Norton.

[44] Vose, M.D. (1991). Generalizing the Notion of Schema in Genetic Algorithms. *Artificial Intelligence, 50,* 385-396.

[45] Vriend, N.J. (1994). *Self-Organized Markets in a Decentralized Economy* (Working Paper 94-03-013) Santa Fe Institute.

[46] Watkins, C.J.C. (1989). *Learning with Delayed Rewards* (Ph.D. thesis) Psychology Department, Cambridge University.

[47] Werbos, P.J. (1974). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences* (Ph.D. thesis) Dept. Applied Mathematics, Harvard University.

[48] Werbos, P.J. (1988). Generalization of Backpropagation with Application to a Recurrent Gas Market Model. *Neural Networks, 1,* 339-356.

[49] Young, H.P. (1991). *Conventional Equilibria* (mimeo) University of Maryland.