

Towards a methodology for experiments with autonomous agents

Luis Antunes and Helder Coelho

Faculdade de Ciências, Universidade de Lisboa
Campo Grande, 1749-016 Lisboa, Portugal
Phone: +351-21-7500526, Fax: +351-21-7500084
{xarax, hcoelho}@di.fc.ul.pt

Abstract. We argue that experimental methodologies are harder to apply when self-motivated agents are involved, especially when the issue of decision gains its due relevance in their model. We use a choice-oriented agent architecture to illustrate a means of bridging the distance between the observer and the actors of an experiment. Traditional experimentation has to give way to exploratory simulation, to bring insights into the design issues, not only of the agents, but of the experiment as well. The role of its designer cannot be ignored, at the risk of achieving only obvious, predictable conclusions. We propose to bring the designer *into* the experiment. To accomplish that, we provide a value-based model of choice to represent the preferences of both entities. This model includes mechanisms that allow for explicit bonds between observer and observed. We use the findings of extensive experimentation with this model to compare current experimental methodologies in what concerns evaluation itself.

- Keywords: Experimental methodologies; autonomous agents.
- Conference topics: AI Foundations and Knowledge Representation; Multi-Agent Systems and Distributed AI.
- Submission to: Paper track
- Multiple submissions: yes

1 Context

“Artificial intelligence [is] the problem of designing agents that *do the right thing*.” [13, page 2, original italics]

Agents can be seen as unwanting actors, but gain additional technological interest and use when they have their own motivations, and are left for autonomous labour. But no-one is completely assured that a program does the “right thing,” or all faulty behaviours are absent. If agents are to be used by someone, *trust* is *the* key issue. But, how can we trust an agent that pursues its own agenda to accomplish some goals of ours [3]?

Autonomy deals with the agents’ freedom of choice, and choice leads to the agents’ behaviour through specific phases in the decision process. Unlike BDI (beliefs-desires-intentions) models, where the stress is given on the technical issues dealing with the agents pro-attitudes (what can be achieved, how can it be done), in BVG (beliefs-values-goals) multi-dimensional models, the emphasis is given on choice machinery, through explicit preferences. Choice is about *which* goals to pursue (or, where do the goals come from), and how the agent *prefers* to pursue them (or, which options the agent wants to pick).

The central question is the evaluation of the quality of decision. If the agent aims at optimising this measure (which may be multi-dimensional), why does s/he not use it for the decision in the first place? And, should this measure be unidimensional, does it amount to a utility function (which would configure the “totalitarian” view: maximising expected utility as the *sole* motivation of the agent)? This view, however discredited since the times of the foundation of artificial intelligence [14], still prevails in many approaches, even through economics or the social sciences (cf. [8]).

In this paper we readdress the issue of principled experimentation involving self-motivated agents. The sense of discomfort borne by reductionist approaches undermines the conclusions of the (ever-so-few) experiments carried out in the field. Hence, we issue a contribution for the synthesis of a method for systemic experimental integration.

In the next section, we summarise the choice framework we adopt, and state the problem of evaluating the results of the agents’ decisions. In section 3, we briefly compare two experimental methodologies. We conclude that none completely solves the issue, and note the similarities between evaluation of the results by the designer, and adaptation by the agents. In section 4 we propose two answers for the issue of assessing experimental results. The combination of both approaches, bringing the designer’s insights and conjectures into the setting of experiments, fits well into the notion of pursuing exploratory simulation. In the last two sections, we briefly present some experimental results, and finally conclude by exalting the advantages of explicitly connecting the experimenter’s and the agents’ evaluative dimensions.

2 Choice and evaluation

The role of value as a mental attitude towards decision is to provide a reference framework to represent agent’s preference during deliberation (the pondering of options candidate to contribute to a selected goal). In the BVG choice framework, the agent’s system of values evolves as a consequence of the agent’s assessment of the results of previous decisions. Decisions are evaluated against certain dimensions (that could be the same previously used for the decision or not), and this assessment is fed back into

the agent’s mind, by adapting the mechanisms associated with choice. This is another point that escapes the traditional utilitarian view, where the world (and so the agent) is static and known. BVG agents can adapt to an environment where everything changes, including the agent’s own preferences (for instance as a result of interactions). This is especially important in a multi-agent environment, since the agents are autonomous, and so potentially sources of change and novelty.

The evaluation of the results of our evaluations becomes a central issue, and this question directly points to the difficulties in assessing the results of experiments. We would need meta-values to evaluate those results. But if those “higher values” exist (and so they are the important ones) why not use them for decision?

When tackling the issue of choice, the formulation of hypotheses and experimental predictions becomes delicate. If the designer tells the agent how to choose, how can he not know exactly how the agent will choose? To formulate experimental predictions and then evaluate to what extent they are fulfilled becomes a spurious game: it amounts to perform calculations about knowledge and reasons, and not to judge to what extent those reasons are the best reasons, and correctly generate the choices. We return to technical reasons for behaviour, in detriment of the will and the preferences of the agent.

By situating the agent in an environment with other agents, autonomy becomes a key ingredient, to be used with care and balance. The duality of value sets becomes a necessity, as agents cannot access values at the macro level, made judiciously coincide with the designer values. The answer is the designer, and the problem is methodological. The update mechanism provides a way to put to test this liaison between agent and designer.

The designer’s model of choice cannot be the model of perfect choice against which the whole world is to be evaluated. It is our strong conviction that the perfect choice does not exist. It is a model of choice to be compared to another one, by using criteria that in turn may not be perfect.

3 Experimental Methodologies

When Herbert Simon received his Turing award, back in 1973, he felt the need to postulate “artificial intelligence is an empirical science.” The duality science/engineering was always a mark of artificial intelligence, so that claim is neither empty nor innocent. Since that time, there has been an ever-increasing effort in artificial intelligence and computer science to experimentally validate the proclaimed results.

3.1 A methodology for principled experimentation

Cohen’s MAD (modelling, analysis and design) methodology [6] is further expanded in [7], where he states the fundamental question to link this methodology to the concept of experiment with self-motivated agents: “What are the criteria of good performance? Who defines these criteria?”

The answer to these questions is an invitation to consider rationality itself, and its criteria. The fact that rationality is situated most times imposes the adoption of *ad hoc* decision criteria. But the evaluation of the results of experiments is not intrinsically different from the evaluation the agents conduct of their own performance (and upon which they base their adaptation). In particular, there was always a designer defining both types of evaluation. So the question comes natural: why would the design of some component be “better” than the other (and support one “right thing”)? Most times

there is no reason at all, and the designer uses the same criteria (the same “rationality”) either for the agent’s adaptation or for the evaluation of its performance.

3.2 A methodology from the social sciences

Computational simulation is methodologically appropriate when a social phenomenon is not directly accessible [11]. A new methodology can be synthesised, and designated “exploratory simulation” [8]. The prescriptive character (exploration) cannot be simplistically reduced to optimisation, such as the descriptive character is not a simple reproduction of the real social phenomena.

A recent methodology for computational simulation is the one proposed by Gilbert [10]. This is not far from MAD, but there are fundamental differences: in MAD there is no return to the original phenomenon, the emphasis is still on the system, and the confrontation of the model with reality is done once and for all, and represented by causal relations. All the validation is done at the level of the model, and the journey back to reality is done already in generalisation. In some way, that difference is acceptable, since the object of the disciplines is also different. But it is Cohen himself who asks for more realism in experimentation, and his methodology fails in that involvement with reality.

But, Is it possible to do better? Is the validation step in Gilbert’s methodology a realist one? Or can we only compare models with other models and never with reality? If our computational model produces results that are adequate to what is known about the real phenomenon, can we say that our model is validated, or does that depend on the source of knowledge about that phenomenon? Isn’t that knowledge obtained also from models? For instance, from results of questionnaires filled by a representative sample of the population – where is the real phenomenon here? Which of the models is then the correct one?

The answer could be in [15]: social sciences have an exploratory purpose, but also a predictive and even prescriptive one. Before we conduct simulations that allow predictions and prescriptions, it is necessary to understand the phenomena, and for that one uses exploratory simulation, the exploration of simulated (small) worlds. But when we do prediction, the real world gives the answer about the validity of the model.

Once collected the results of simulations, they have to be confronted with the phenomenon, for validation. But this confrontation is no more than analysis. With the model of the phenomenon to address and the model of the data to collect, we have again a simplification of the problem, and the question of interpretation returns. It certainly isn’t possible to suppress the role of the researcher, ultimate interpreter of all experiments, be it classical or simulation.

4 Two answers

In this section we will present two different answers for the problem of analysing (and afterwards, generalising) the results of the experimentation, which we have already argued to have quite a strong connection to the problem of *improving* the agents performance as a result of evaluation of the previous choices.

The explicit consideration of the relevant evaluative dimensions in decision situations can arguably provide a bridge between the agent’s and the experiments designer’s mind. In a multi-dimensional choice model, the agent’s choice mechanisms are fed back with a set of multi-dimensional update values. These dimensions may or not be the same that were used to make the decision in the first place. If these dimensions should

be different, we can identify the ones that were used for decision with the interests of the agent, and the ones used for update with the interests of the designer. And moreover, we have an explicit link between the two sets of interests. So, the designer is no longer left for purely subjective guessing of what might be happening, confronted with the infinite regress of ever more challenging choices. S/he can explore the liaisons provided by this choice framework, and experiment with different sets of preferences (desired results), both of hers and of the agents.

4.1 Positivism: means-ends analysis in a layered mind

We can postulate a positivist (optimistic) position by basing our ultimate evaluations on a pre-conceived ontology of such deemed relevant dimensions (or values). Having those as a top-level reference, the designer's efforts can concentrate on the appropriate models, techniques and mechanisms to achieve the best possible performance as measured along those dimensions.

It seems that all that remains is then optimisation along the desired dimensions, but even in that restrained view we have to acknowledge that it does not mean that all problems are now solved. Chess is a domain where information is perfect and the number of possibilities is limited, and even so it was not (will it ever be?) solved.

Alternatively, the designer can be interested in evaluating how the agents perform in the absence of the knowledge of what dimensions are to be optimised. In this case, several models can be used, and the links to the designer's mind can still be expressed in the terms described above.

The key idea is to approximate the states that the agent wishes to achieve to those that it believes are currently valid. This amounts to performing a complex form of means-ends analysis, in which the agent's sociality is an issue, but necessarily one in which the agent does not have any perception about the meta-values involved. Because that would reinstate the infinite regression problem.

The external evaluation problem can be represented in terms as complex as the experiment designer thinks appropriate. In a BDI-like logical approach, evaluation can be as simple as answering the question "were the desired states achieved or not?," or as complicated as the designer desires and the decision framework allows to represent. The choice mechanisms update becomes an important issue, for they are trusted to generate the desired approximation between the agent's performance (in whichever terms) and the desired one.

Interesting new architectural features recently introduced by Castelfranchi [4] can come to the aid of the task of unveiling these ultimate aims that justify behaviour. Castelfranchi acknowledges a problem for the theory of cognitive agents: "how to *reconcile the 'external' teleology of behaviour with the 'internal' teleology governing it*; how to reconcile intentionality, deliberation, and planning with playing social functions and contributing to the social order." [4, page 6, original italics].

Castelfranchi defends *reinforcement* as a kind of internal natural selection, the selection of an item (e.g. a habit) directly within the entity, through the operation of some internal choice criterion. And so, Castelfranchi proposes the notion of *learning*, in particular, reinforcement learning in cognitive, deliberative agents. This could be realised in a hybrid layered architecture, but not one where reactive behaviours compete against a declarative component. The idea is to have "a number of low-level (automatic, reactive, merely associative) mechanisms operate *upon* the layer of high cognitive representations" [4, page 22, original italics].

Damasio's [9] somatic markers, and consequent mental reactions of attraction or repulsion, serve to constrain high level explicit mental representations. This mental

architecture can do without the necessity of an infinite recursion of meta-levels, goals and meta-goals, decisions about preferences and decisions. In this meta-level layer there could be no explicit goals, but only simple procedures, functionally teleological automatisms.

In the context of our ontology of values, the notion of attraction/repulse could correspond to the top level of the hierarchy, that is, the ultimate value to satisfy. Optimisation of some function, manipulation and elaboration of symbolic representations (such as goals), pre-programmed (functional) reactivity to stimuli, are three faces of the same notion of ending up the regress of motivations (and so of evaluations over experiments). This regress of abstract motivations can only be stopped by grounding the ultimate reason for choice in concrete concepts, coming from *embodied* minds.

4.2 Relativism: extended MAD, exploratory simulation

There are some problems in the application of MAD methodology to decision situations. MAD is heavily based on hypotheses formulation and predictions about systems behaviour, and posterior confrontation with experimental observations. An alternative could be conjectures-led exploratory simulation.

The issues raised by the application of MAD deal with meta-evaluation of behaviours (and so, of underlying models). We have proposed an extension to MAD that concerns correction between the diverse levels of specification (from informal descriptions to implemented systems, passing by intermediate levels of more or less formal specification). This extension is based on the realisation of the double role of the observer of a situation (which we could translate here into the role of the agent and that of the designer).

The central point is to evaluate the results of agent's decisions. Since the agent is autonomous and has its own reasons for behaviour, how can the designer dispute its choices? A possible answer is that the designer is not interested in allowing the agent to use the best set of reasons. In this case what is being tested is not the agent, but what the designer thinks are the best reasons. The choice model to be tested is not the one of the agent, and the consequences may be dramatic in open societies.

In BVG, the feedback of such evaluative information can be explicitly used to alter the agents choice model, but also to model the mind of the designer. So, agents and designer can share the same terms in which the preferences can be expressed, and this eases up validation. The model of choice is not the perfect reference against which the world must be evaluated (such a model cannot exist), but just a model to be compared to another one, by using criteria that again might not be perfect.

This seems to amount to an infinite regress. If we provide a choice model of some designer, it is surely possible to replicate it in the choice model of an agent, given enough liberty degrees to allow the update mechanisms to act. But what does that tell us? Nothing we couldn't predict from the first instant, since it would suffice that the designer's model would be used *in* the agent. In truth, to establish a realist experiment, the designer's choice model would itself be subject to continuous evolution to represent his/her choices (since it is immersed in a complex dynamical world). And the agent's model, with its update mechanisms, would be "following" the other, as well as it could. But then, what about the designer's model, what does it evolve to follow? Which other choice model can this model be emulating, and how can it be represented?

Evaluation is harder for choice, for a number of reasons: choice is always situated and individual, and it is not prone to generalisations; it is not possible to establish criteria to compare choices that do not challenge the choice criteria themselves; the

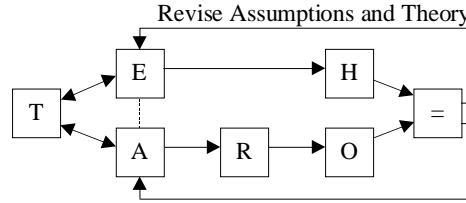


Fig. 1. Construction of theories. An existing theory (T) is translated in a set of assumptions (A) represented by a program and an explanation (E) that expresses the theory in terms of the program. The generation of hypotheses (H) from (E) and the comparison with observations (O) of runs (R) of the program allows both (A) and (E) to be revised. If finally (H) and (O) correspond, then (A), (E) and (H) can be fed back into a new revised theory (T) that can be applied to a real target (from [11]).

adaptation of the choice mechanisms to an evaluation criteria appears not as a test to its adaptation capabilities, but rather as a direct confrontation of the choices.

Who should tell if our choices are good or not, based on which criteria can s/he do it, why would we accept those criteria, and if we accept them and start making choices by them, how can we evaluate them afterwards? By transposing this argument to experimental methodology, we see the difficulty in its application, for the decisive step is compromised by this opposition between triviality (when we use the same criteria to choose and to evaluate choices) and infinite and inevitable regression (that we have just described).

Despite all this, the agent *cannot* be impotent, prevented from improving its choices. Certainly, human agents are not, since they keep choosing better (but not every time), learn from their mistakes, have better and better performances, not only in terms of some external opinion, but also according to their own. As a step forward, and out of this uncomfortable situation, we can also consider that the agent has two different rationalities, one for choice, another for its evaluation and subsequent improvement. One possible reason for such a design could be the complexity of the improvement function be so demanding that its use for common choices would not be justified.

To inform this choice evaluation function, we can envisage three candidates: (i) a higher value, or some specialist's opinion, be it (ii) some individual, or (iii) some aggregate, representing a prototype or group.

The first, we have already described in detail in the previous subsection: some higher value, at a top position in an ontological hierarchy of value. In a context of social games of life and death, survival could be a good candidate for such a value. As would some more abstract dimension of goodness or righteousness of a decision. That is, the unjustifiable (or irreducible) sensation that, all added up, the right (good, just) option is evident to the decider, even if all calculations show otherwise. This position is close to that of *moral imperative*, or *duty*. But this debate over whether all decisions must come from the agents pursuing their own interest has to be left for further studies.

The second follows Simon's idea for the evaluation of choice models: choices are compared to those made by a human specialist. While we want to verify if choices are the same or not, this idea seems easy to implement. But if we want to argue that the artificial model chooses *better* than the reference human, we return to the problem of deciding what 'better' means.

The third candidate is some measure obtained from an aggregation of agents which are similar to the agent or behaviour we want to study. We so want to compare choices

made by an agent based on some model, with choices made by some group to be studied (empirically, in principle). In this way we test realistic applications of the model, but assuming the principle that the decider agent represents in some way the group to be studied.

4.3 Combining the two approaches

A recent methodological approach can help us out here [12]. The phases of construction of theories are depicted in figure 1. However, we envisage several problems in the application of this methodology. Up front, the obvious difficulties in the translation from (T) to (E) and from (T) to (A), the subjectivity in the selection of the set of results (R) and corresponding observations (O), the formulation of hypotheses (H) from (E) (as Einstein said: “no path leads from the experience to the theory”). The site of the experimenter becomes again central, which only reinforces the need of defining common ground between him/her and the mental content of the agents in the simulation.

Thereafter, the picture (as its congeners in [12]) gives further emphasis to the traditional forms of experimentation. But Hales himself admits experimentation in artificial societies demands for new methods, different from traditional induction and deduction. Like Axelrod says: “Simulation is a third form of making science. (...) While induction can be used to discover patterns in data, and deduction can be used to find consequences of assumptions, the modelling of simulations can be used as an aid to intuition.” [2, page 24]

This is the line of reasoning already defended in [8]: to observe theoretical models running in an experimentation test bed, it is ‘exploratory simulation.’ The difficulties in concretising the verification process (=) in figure 1 are even more stressed in [5]: the goal of these simulation models is not to make predictions, but to obtain more knowledge and insight.

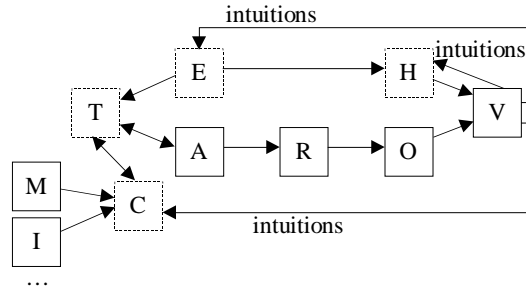


Fig. 2. Exploratory simulation. A theory (T) is being built from a set of conjectures (C), and in terms of the explanations (E) that it can generate, and hypotheses (H) it can produce. Conjectures (C) come out of the current state of the theory (T), and also out of metaphors (M) and intuitions (I) used by the designer. Results (V) of evaluating observations (O) of runs (R) of the program that represents assumptions (A) are used to generate new explanations (E), which allow the reformulation of the theory (T).

This amounts to radically changing the drawing of figure 1. The theory is not necessarily the starting point, and the construction of explanations can be made autonomously, as well as the formulation of hypotheses. Both can even result from the application of the model, instead of being used for its evaluation. According to Casti

[5], model validation is done qualitatively, recurring to intuitions of human specialists. These can seldom predict what occurs in simulations, but they are experts at explaining the occurrences.

Figure 2 is inspired in the scheme of explanation discovery of [12], and results from the synthesis of the scheme for construction of theories of figure 1, and a model of simulations validation. The whole picture should be read at the light of [5], that is, the role of the experimenter and his/her intuition is ineluctable. Issues of translation, retroversion and their validation are important, and involve the experimenter. On the other hand, Hales' (=) is substituted by an evaluation machinery (V), that can be designed around values. Here, the link between agents and experimenter can be enhanced by BVG choice framework.

One of the key points of the difference between figures 1 and 2 is the fact that theories, explanations and hypotheses are being constructed, and not only given and tested. Simulation is precisely the search for theories and hypotheses. These come from conjectures, through metaphors, intuitions, etc. Even evaluation needs intuitions from the designer to lead to new hypotheses and explanations. This process allows the agent's choices to approximate the model that is provided as reference. Perhaps this model is not as accurate as it should be, but it can always be replaced by another, and the whole process of simulation can provide insights into what this other model should be.

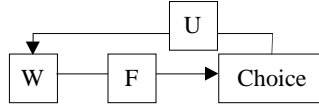


Fig. 3. Choice and update in the BVG architecture.

The move from BDI to BVG was driven by a concern with choice. But to tune up the architecture, experimentation is called for. BVG is more adaptive to dynamic situations than BDI, and this presents new demands on the experimental methodology. In BVG (see figure 3), choice is based on the agent's values (W), and performed by a function F . F returns a real value that momentarily serialises the alternatives at the time of decision. The agent's system of values is updated by a function U that uses multidimensional assessments of the results of previous decisions. We can represent the designer's choice model by taking these latter dimensions as a new set of values, W' . Mechanisms F and U provide explicit means for drawing the link between the agent's (choosing) mind and the designer's experimental questions, thus transporting the designer into the (terms of the) experiment. This is accomplished by relating the backwards arrows in both figures (2 and 3). We superimpose the scheme of the agent on the scheme of the experiment.

5 Assessment of experimental results

This concern with experimental validation is an important keynote in the BVG architecture. Initially we reproduced (using Swarm) the results of Axelrod's "model of tributes," because of the simplicity of the underlying decision model [1]. Through principled exploration of the decision issues, we uncovered certain previously unidentified features of the model. But the rather rigid character of the decision problems would not allow the model to show its full worth.

In other experiments, agents selected from a pool of options, in order to satisfy some (value-characterised) goals. This introduced new issues in the architecture, such as non-transitivity in choice, the adoption of goals and of values, non-linear adaptation, the confront between adaptation based on one or multiple evaluations of the consequences of decisions. We provide some hints into the most interesting results we have found.

In a series of runs, we included in F a component that subverts transitivity in the choice function: the same option can rise different expectations (and decisions) in different agents. A new value was incorporated, to account for the effect of surprise that a particular value can raise, causing different evaluations (of attraction and of repulse).

The perils of subverting transitivity are serious. It amounts to withdrawing the golden rule of classical utility, that “all else being equal” we will prefer the better option. However, we sustain that it is not necessarily irrational (sometimes) not to do so. We have all done that in some circumstances. The results of the simulations concerning this effect of surprise were very encouraging. Moreover, the agent’s choices remained stable with this interference. The agent does not loose sense of what its preferences are, and what its rationality determines. It acts as if it allowed itself a break, in personal indulgence.

In other runs, we explored the role of values in regulating agent interactions, for instance, goal adoption. We found that when we increase the heterogeneity of the population in terms of values (of opposite sign, say), we note changes in the choices made, but neither radical, neither significant, and this is a surprising and interesting fact. The explanation is the “normalising” force of the multiple values and their diffusion. An agent with one or another different value still remains in the same world, sharing the same information, exchanging goals with the same agents. The social ends up imposing itself.

What is even more surprising is that this force is not so overwhelming that all agents would have exactly the same preferences. So many things are alike in the several agents, that only the richness of the model of decision, allied to their particular life stories, avoids that phenomenon.

The model of decision based on multiple values, with complex update rules, and rules for information exchange and goal adoption, presents a good support for decision making in a complex and dynamic world. It allows for a rich range of behaviours that escapes from directed and excessive optimisation (in terms of utilitarian rationality, it allows for “bad” decisions), but does not degenerate in pure randomness, or non-sense (irrationality). It also permits diversity of attitudes in the several agents, and adaptation of choices to a dynamic reality, and with (un)known information.

6 Conclusions

No prescribed methodology will ever be perfect for all situations. Our aim here is to draw attention to the role of the designer in any experiment, and also to the usually underaddressed issue of choice in the agent’s architecture. Having a value-based choice model at our hands as a means to consider self-motivated autonomous agents, these two ideas add up to provide a complete decision framework, where the designer is brought into the experiment, through the use of common terms with the deciding agents. This is a step away from reductionism, and towards a holistic attitude in agent experimentation.

References

1. Robert Axelrod. A model of the emergence of new political actors. In *Artificial Societies – The Computer Simulation of Social Life*. UCL Press, 1995.
2. Robert Axelrod. Advancing the art of simulation in the social sciences. In *Simulating Social Phenomena*, volume 456 of *LNEMS*. Springer, 1997.
3. Cristiano Castelfranchi. Guarantees for autonomy in cognitive agent architecture. In *Intelligent Agents: agent theories, architectures, and languages, Proc. of ATAL'94*, volume 890 of *LNAI*. Springer, 1995.
4. Cristiano Castelfranchi. The theory of social functions: challenges for computational social science and multi-agent learning. *Journal of Cognitive Systems Research*, 2, 2001.
5. John L. Casti. Would-be business worlds. *Complexity*, 6(2), 2001.
6. Paul R. Cohen. A survey of the eighth national conf. on AI: Pulling together or pulling apart? *AI Magazine*, 12(1):16–41, 1991.
7. Paul R. Cohen. *Empirical Methods for AI*. The MIT Press, 1995.
8. Rosaria Conte and Nigel Gilbert. Introduction: computer simulation for social theory. In *Artificial Societies: the computer simulation of social life*. UCL Press, 1995.
9. António Damásio. *Descartes' error*. Putnam's sons, New York, 1994.
10. Nigel Gilbert. Models, processes and algorithms: Towards a simulation toolkit. In *Tools and Techniques for Social Science Simulation*. Physica-Verlag, 2000.
11. Nigel Gilbert and Jim Doran, editors. *Simulating Societies: the computer simulation of social phenomena*. UCL Press, London, 1994.
12. David Hales. *Tag Based Co-operation in Artificial Societies*. PhD thesis, Univ. Essex, 2001.
13. Stuart Russell and Eric Wefald. *Do the right thing – studies in limited rationality*. The MIT Press, 1991.
14. Herbert A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, Feb. 1955.
15. Klaus G. Troitzsch. Social science simulation – origins, prospects, purposes. In *Simulating Social Phenomena*, volume 456 of *LNEMS*. Springer, 1997.