

The Practical Modelling of Context-Dependent Causal Processes – A Recasting of Robert Rosen’s Thought

Bruce Edmonds
Centre for Policy Modelling
Manchester Metropolitan University

Abstract

The paper is a critical recasting of some of Robert Rosen’s thought. It is argued that a lot of the thrust of his work can be better understood when re-cast in terms of the context-dependency of causal models. When re-cast in this way I seek to highlight how his thought does not lead to the abandonment of formal modelling and a descent into relativism but a more careful and rigorous science of complex systems. This also sheds light on several aspects of modelling including: the need for multiple models; the nature of modelling noise; and why adaptive systems cause particular problems to modellers. In this way I hope to decrease researchers’ fear that, by taking Rosen’s criticisms seriously, they would have to abandon the realm of acceptable science.

Introduction

Robert Rosen was a pioneer. He identified and raised a number of important issues to do with modelling complex phenomena. In particular he: examined and analysed the process of scientific modelling itself [1]; he pointed out the impossibility of the idealised reductionist and universalist project [2]; he pointed out the special nature of systems that anticipated the future [3]; and he drew a sharp distinction between complex and simple systems [4]. He was a classic prophet in that he correctly anticipated present trends and difficulties – he was a voice that was sometimes heard and occasionally understood, but not accepted into the mainstream because it would have involved too radical a shift for most scientists. I think people were worried that accepting Rosen’s critique would have involved the abandonment of formal modelling and the acceptance of relativism, which would put them beyond the realm of acceptable science. I think that they were mistaken.

In this paper I seek to recast these issues and ideas in a new light and argue that neither the abandonment of formal modelling nor a descent into relativism are necessary consequences of this shift. On the contrary, I would like to show that accepting this shift can lead to better and more careful science. The main crux of my argument is that causation is inherently context-dependent but that context-dependency is not the same as relativism. From these central insights comes new ways of understanding, including: the inevitability of multiple models; the nature of noise; and some areas where we would expect modelling to be inherently context-dependent.

In this it will become clear that, although appreciative of his analyses and fully acknowledging his pioneering work (which I build upon), I disagree with him in some respects. To take two examples: I think that the division between what is complex and simple is more complicated than in his accounts; and I think that any attempt to conclusively demonstrate his (in my view correct) critiques of the accepted scientific

world view are bound to fail – how we *should* approach the modelling of complex phenomena is not a matter of proof, but more concerned with practically grounded issues and constraints [1].

This paper is structured as follows. I start by briefly discussing the ideas of causation and context, before moving on to argue that causation is an inherently context-dependent idea. I then point out that context-dependency is not the same as subjectivity and that this is compatible with the project of careful science. This picture then has some implications which are then briefly touched upon in the following sections: the nature of modelling noise; the appearance of multiple models when dealing with complex phenomena; Rosen's distinction between simple and complex systems; and the special case of systems that are involved in some sort of modelling or adaptation themselves.

Any paper which seriously wishes to build upon the work of Rosen has to face up to foundational issues to do with the nature of science itself – this is such a paper. Hence it is important that I try to avoid “building in” assumptions concerning its foundations that I wish to discuss. For this reason I will avoid the use of certain terms, notably “system” (as in “the lymphatic system”), replacing it with various circumlocutions such as “phenomena”, “observations” and the like. This is not to indicate any adherence to a “post-modernist” style but because it seeks to analyse the assumptions that are already assumed by the use of such terms. Indeed the whole tenor of this paper is that an examination of the foundations of science (understanding and thinking about what one is doing in science and why, as well as just doing it) does not involve a shift towards post-modernism and relativism but can lead towards a greater level of scientific rigour.

Causation

If there were a technique that, given a correct description of some class of phenomena, correctly predicted what would happen, but where there was absolutely *no* understanding of how it worked then this would not count as science. I will call such hypothetical technique a “magic technique”. It is likely that *if* any such magic technique *were* discovered then a huge amount of effort would be expended in finding out what the technique did and why it predicted that class of phenomena. Indeed, it is the faith that there *are no such magic techniques* that underlies the scientific enterprise – a belief that if we can predict then we can also (eventually) explain. What is missing from the magic technique is a convincing account of the causation within the technique (or as is more usually the case: chains of inference) from initial conditions to its results and how that account relates to the causation within the target phenomena. It is the mapping between the causation in the model and that in the phenomena that makes it scientific [6]. Thus causation is at the heart of the scientific enterprise.

The converse situation where, there is a “good” account of the causation acting within some class of phenomena but where there is no good technique for predicting the outcomes, is quite common in science. It is one of the hallmarks of *complex* phenomena that we may understand it (in some sense) but not be able to predict it in all its aspects even if we had a complete and precise description of all its relevant aspects. Unlike the previous case of “magic techniques” such complex phenomena

are the norm in science. Where Rosen differed from many of his colleagues is his denial of their faith that all such understandable phenomena (including those that are, in some sense, complex) are capable of being brought within a single, formal and predictive framework. Rosen's denial of this faith was an, *in principle* denial – he spent much of his writing seeking to demonstrate that, even if the severe practical difficulties were to be overcome in terms of measurement, computation etc. that there will be many cases where this faith was wrong. Whilst I would guess that Rosen was fundamentally right in this respect, I do not think that this is demonstrable [1]. If I am right it will always be an open question whether any class of phenomena is amenable to being adequately formalised by a single predictive model or not. Those with the faith may continue in their quest; those without, looking for other ways forward or other problems to tackle [1].

The nature of causation has been hotly disputed within philosophy from Aristotle onwards, with some disputing whether the term has any coherent and useful meaning at all, e.g. Hume. There is not space to go into these disputes here. However, it is notable that this history of doubt and equivocation is in stark contrast to the certainty and confidence that practicing scientists have in using the notion. To them causation is a “bread and butter” reality – the reality of causation is not in doubt, a reality that is continually reinforced by their experience, a reality that they use to guide their investigations.

Why do scientists and philosophers have such different views about causation? The answer lies, I think, in the generality that is attempted within these different fields. Philosophy seeks universal answers – that is, if another philosopher can come up with a valid counter example to any thesis (however weird the example is) then the burden is on the proponent to alter that thesis to deal with it. Thus philosophical accounts of causation attempt to give an account that is completely context-free and this results in either a negative conclusion (there is no such thing as causation) or complicated and qualified accounts. Scientists, on the other hand, work within specific (and usually highly familiar) contexts trying (and in many cases succeeding) to produce causally-based models of what they are observing. For them, causation works. The answer, I suggest, lies in the essential context-dependency of causation.

Context

Context is one of these problematic words (like “complexity”) that combines being fashionable with having a diversity of different uses. Thus although everybody has some idea what these terms mean, a close examination reveals a lot of different possible precise meanings [7]. This is sort of “common vagueness” is often true of those negative concepts (such as context and complexity) that are often defined by what they are not or by some quality that they lack. Thus complexity is often merely what is not simple (in some sense) and something is explained by the context if there is no obvious foreground feature that might indicate it. There is now a considerable literature on context, including an international series of conferences on the subject.

I will introduce context with respect to the modelling process, for more detail about this see [7]. In the natural world there will almost always be an indefinite number of factors that could possibly influence any particular complex real event (if one does not limit the scope to something sensible or what is likely to occur). In order to obtain

a formal model one has to exclude most of these potential factors, implicitly making suitable assumptions such as that this factor will remain constant or that one will not occur [9]. This set of possible assumptions is potentially unlimited (for example person A will not disrupt the experiment; person B will not disrupt the experiment etc.). However, *given* this huge set of background assumptions it can be possible to form crisp, formal models that are useful for predicting aspects of the phenomena of interest. The set of factors that are not in the foreground model but are necessary for that model to be useful implicitly defines the scope of that model. When a set of models are clustered so that their scopes are similar – in other words the scopes cluster – this might indicate a “context”. Thus a context can be associated with a cluster of models and facts that are useful in that context. If every different model had a totally different scope there would not be any natural field or domain in which to seek understanding – no identifiable contexts. Thus, classically we have the experimental context, where many of these potential factors are controlled for by the design of the experiment and the cellular context where defined by what is a particular living cell controls for – corresponding to the classic *in vitro*, *in vivo* distinction.

Thus the basic “context heuristic” is to recognise these clusters and then do crisp and limited reasoning/induction *within* these contexts. There is evidence that many human cognitive abilities are context-dependent, including: memory, language, reasoning and pairwise-judgement [10]. Thus presumably humans use some sort of largely unconscious pattern-recognition neural machinery to learn and distinguish these contexts. This divides up the cognitive load between rich but “fuzzy” learning and recognition processes *of* contexts and precise modelling and reasoning processes *within* contexts. Humans generally perform only limited or simple reasoning (except when diagnosing faults or problem solving) but seem to be good at switching between different ways of considering the same problem. Note that only sometimes is a context amenable to explicit identification as such by such means as introspection etc. Many times we can correctly identify a context but will have difficulty defining it or naming it.

Thus the heuristic of context makes possible the use of relatively simple, crisp models in a complex world. Of, course what is crucial is that these models not only can be useful within a particular context but that they can be used to transmit useful knowledge from one situation to another. This communication of knowledge works under the following conditions (from [7]):

1. that some of the possible factors influencing an outcome are separable in a practical way;
2. that a useful distinction can be made between those factors that can be categorised as foreground features (including 'causes') and the others;
3. that the background factors are capable of being recognised later on;
4. that the world is regular enough for such models to be at all learnable;
5. that the world is regular enough for such learnt models to be at all useful when applied in situations where the context can be recognised.

Condition 3 is crucial in this, the context where a model is valid (for the purposes of this discussion – “valid” means “useful”) needs to be reliably recognisable by others. This recognition can be as a result of formal definition and analysis, but is often more dependent upon a richer but less exact learning process that is built up as the result of experience. Fortunately us humans seem to have such a context-recognition machine in our neural apparatus. Thus we all unfailingly recognise life when we see it, but that does not mean that it is easy to define it precisely (I am not aware of any undisputed definition of life).

The difficulties arise when this condition is not reliable, for example when there are multiple possible contexts discernable from a single situation. Thus if one is examining any particular mechanism or process within a cell, there will be many aspects that this can be considered from. Often these are a result of the different purposes for which the process is being examined.

Here I need to distinguish two different uses of the word “context”. In common usage “context” refers to the external situation where the clustering of models, knowledge, assumptions, labels, etc. apply. However to talk about context-dependency and contexts one needs to generalise over similar contexts, for example over the contexts of model learning and applications talked about above. Thus what we are often really talking about here (and in similar discussions) is not referring to an external situation but the cognitive or modelling correlate of a set of similar situations. This correlate is sometimes called the cognitive or internal context and distinguished from the external context. If the context heuristic is to work it must be that it is possible to learn/develop such correlates such that they adequately relate to the set of external situations they are associated with – otherwise the above heuristic would not work. Thus in the discussion below I am often referring to the internal version of context, because that is the version we have modelling access to, the external context necessarily remains an idealisation and is not directly accessible.

The context-dependency of causation

A cause is a something that is necessary for an event to occur. It can be thought of in the negative as a counterfactual: if that thing did *not* exist the event would *not* have occurred. Of course, what is not specified in this negative reformulation is what has to remain constant whilst the counterfactual of the thing exists or not is conceptually tried out. Thus most formal accounts of causation (e.g. [11]) start with the key assumption that all possible causes have been listed and then produces a formal (graph theoretic) method of distinguishing which of these will count as a cause.

Thus negative assumptions (this is constant, that will not occur) or normality assumptions are necessary in order to make the notion of a cause meaningful (at least to us finite beings). It is only obvious what should be considered normal with respect to a given context, so normality assumptions are equivalent to a host of negative assumptions associated with the context(s) where the normality assumptions hold.

Let me illustrate this with an example. Let us suppose that we are trying to understand how the chemical reactions effect some aspect of a cell that they occur in. Given how interrelated the metabolic pathways in a cell are, if one traced the possible chain of causation back in time then one could identify a huge number of possible causes for any particular event. Formally “the” cause of the event would be of the

form (A and B and C and ...) or (M and N and O and ...) or (W and X and Y and ...) or ..., where the letters stand for (possibly) different causes. It is likely that the further one traced the chains of causation backwards in time the longer and more complicated the expression would be. Such “wild causation” [12] makes the whole notion of causation unusable, and ultimately meaningless – almost everything would be causing everything. However if one is considering the cell from the viewpoint of a particular context (say only those pathways to do with the transport of Sodium) then this would limit the causal spread, opening the possibility that it might be formally modelled. The resulting model of sodium transport might work well for the purpose of predicting the concentration of sodium ions in the cell, under circumstances where the non-sodium chemical reactions in the cell do not effect this, but in other contexts (e.g. when something triggers a new set of reactions) the model may well fail. What *causes* changes in the concentration of sodium ions depends upon what one considering irrelevant, which depends upon the context under which one is seeking understanding. However, the context of the living cell is reliably recognisable by all in the field and the sub-contexts defined by particular interests usually distinguishable (albeit sometimes with some effort).

Context-dependency is not the same as relativism

Now it should be clear why asserting the context-dependency of causation is not the same as relativism. It does appear that the world is amenable to being usefully decomposed into clusters of situations so that the cognitive heuristic of context works. Thus in many cases we find that we *can* reliably recognise the appropriate context and do useful crisp modelling and reasoning within this context. Furthermore the same context (for all practical purposes) can be recognised by different people, even if they have different educations, cultures, etc. (though sometimes some training or other learning process is required). Thus the truth of crisp models that presume a context does not depend upon the point of view of the person or which person is judging the veracity – one can not make anything true by changing the way one considers things. Valid science, even of complex phenomena, is possible!

The heuristic of dividing the world up into separate contexts and reasoning within these contexts works is not a *necessary* fact but a *contingent* one on the way the part of the universe we inhabit and observe happens to be (including critically, us). It could have been that the universe was so intricately interrelated that this trick did not work – that it was possible that if one did not include *everything* in one’s model that it would be useless! Of course, if this was the case it is difficult to imagine how life could have evolved; how any organism could have gained sufficient “purchase” over its environment that it could have learned to exploit it and support its own reproduction. Thus the only “necessity” for the amenability of our environment to modelling comes from our own structure. Thus this is a sort of anthropomorphic necessity – it might be that (our corner of) the universe “has” to be like this because it has supported the evolution of beings like us able to partially understand it so [13]. It might be perfectly possible that the universe was such that we were not able to observe and think about it in this way, but then we would probably not be here to do so. Such a universe is not one we can talk about (or are talking about)!

However, an awareness of the fact that our models are context-dependent can help alert us to when the application of any particular model might be dubious. This could

be when a model is being applied in a very different context to that it was developed for, or where the system being studied is so complex that people may be studying the same things but different aspects of those things. The point is that one can not just *assume* that a model that is valid in one context is going to be valid in a new context.

The nature of modelling noise

One of the benefits of this way of looking at things is that it also gives us a coherent account of what “modelling noise” is. Given that our models depend upon the way the complex world partially “divides” into contexts, our models will be also. However the division is never perfect so that there are always influences that “leak” between contexts making the results of any within context model less than perfect and predicting the in-context phenomena. This extra-contextual leakage is often characterised (by analogy with electrical phenomena) as “noise” [14]. This noise can not always be brought into the model as an endogenous process (see the section immediately below), in which case it is not predictable by processes within the model. We often use a pseudo-random process to stand-in for this noise on the grounds that this is reasonable for an essentially unpredictable or irrelevant process.

However it should be clear that such extra-contextual influences that are not thought to be crucial for the modelling process, are not necessary random in that they might not have well defined or stationary probability distributions. Thus assuming that any such “noise” is random might not be a good idea – for example in [15] it shown that what appears as noise does not scale with the law of large numbers as the system increases in size. If it were random noise the variance would decrease as a proportion of system size (in bigger systems, the “signal” stands out from the noise), whilst this is patently not the case in some complex systems [16].

Simple and complex phenomena

Rosen sought to draw a clean distinction between simple and complex phenomena by identifying these with formal and natural processes respectively [4]. Thus a machine, in Rosen's scheme, is not anything physical but an idealisation of something physical which is then conflated with the physical things it idealises. Thus a car is constructed so that it behaves, in most contexts we encounter it, very like its formal model would predict. Traditionally science has concerned itself with such “machines”, even in biology, in the sense that when it did encounter complex phenomena (such as organisms) it only focused on those aspects of them that could be thought of as machine-*like*.

Rosen's criterion for a simple system was one that could be modelled by a single encompassing formal model. That is all aspects of the target phenomena that were relevant were predictable (in principle) using a formally defined system of inference. Within the terms introduced here, this would necessitate that all the contexts the complex phenomena were modelled under (and the causation within then) being combinable into a single mega-context (i.e. de-contextualised). Rosen argued that this ruled out *any* natural system being simple and hence identified the simple/complex dichotomy with the formal/natural one. The trouble with this approach is that it denies any complexity comparison between natural systems – they are all equally, ultimately complex. Thus, according to Rosen's picture, the world before life

appeared was just as complex as after it had substantially evolved. I think that what Rosen was concerned about was to attack the unthinking *assumption* that scientific models have universal scope – it was an essentially anti-reductionism and anti-universalist argument. However he sought to demonstrate these arguments in a universal way, by general philosophical demonstration. Adding in context makes sense of Rosen’s approach, *as well* as that of common-sense complexity comparisons and the like.

If we can accept that all modelling is *within* a context (or contexts) then we can accept the thrust of Rosen’s point that a model will never capture all aspects of any phenomena but also maintain that complexity comparisons *within a context* can be meaningful – for example that the production of a new catalyst within a bio-chemical system can change the complexity of that system.

Likewise I would suggest that (unlike Rosen *on the whole*) that formal models can be meaningfully said to be complex in Rosen’s sense. To give a demonstrable example (although most people do not need a demonstration to believe that parts of formal mathematics can be truly complex), consider the case of formal computability (as defined by Turing [17] and others. The “limited halting problem”, H_n , of deciding whether a program of length less than a parameter, n , will eventually halt given an input of less than n , *is computable* given any particular fixed n (it could be implemented as a huge 2D look-up table with rows for program numbers, columns for inputs and true/false as entries). However it is not possible to construct a program that will take a third variable input to specify n (this would solve the famous Turing Halting program which he showed was impossible). In other words all the H_n can not be combined into a single, general program H . In other words, program execution is complex by Rosen’s definition. This arguments is dealt with in more depth in [18].

I think Rosen, proposed the identification of the simple/complex with the formal/natural for two reasons: to highlight the distinction between the thing-itself and the model it is conflated with (which will only be reliable in context); and as a polemical device to attack “complexity science” which, from his point of view, was as reductionist as any other kind of science. However, if I am right the division between the simple and the complex is complicated and dependent upon the context *the division itself* is being considered in.

Multiple models of phenomena

Given that the same complex phenomena might be considered from different viewpoints (different aspects being considered as relevant in each case), where it is impractical to combine them all into one model (leaving aside the arguments of Rosen why this might be impossible *in principle*), then we might well have a situation where we have multiple but different models of the same focus phenomena. This seems to be an inevitable consequence of dealing with complex phenomena, that clusters of related models might result from the study of something rather than a single model/answer [19]. This is in contrast to the simple case where something can be adequately modelled from within a single context.

The trouble is that science has not yet developed an adequate and systematic methodology to deal with such clusters, or at least has not codified and scrutinised such (how to do this sort of thing is mostly imparted *implicitly* during the training of a

scientist by example). At the moment there are three accepted ways: to show that one model is a special case of another (*specialisation*); to derive a simpler model from another using suitable approximations (*approximation*); and to show that two differently constructed models come up with similar (or different) results (*comparison*). Specialisation is only applicable for simple systems, i.e. within a single context or set of assumptions. Comparison requires that the results are comparable, at least within some domain – this does not mean that they have to be based on the same assumptions and could each be dealing with very different aspects, but there does have to be some overlap in terms of scope so the relevant contexts can not be too dissimilar. The best that can be hoped for in this case is an approximate consistency with the intersection of the models' scopes. Approximation is an interesting case: approximations usually rely upon additional assumptions that come from knowledge about the context under which the target phenomena is being considered. Thus different approximations to some general formulation lead to models within different contexts. Cartwright [20] reports on the use of parallel but different approximations which broadly agree on the same predictions might give a scientist confidence in the results.

These three “moves” are not the only possible ones. For example it may models are iterated, in that one model is modelling another, like when an individual-based computational model results in a equation which, in turn, models a set of data. However little is known about other ways of combining and maintaining clusters of models, except by discovered rules of thumb. This area cries out for development.

Anticipatory systems

A special difficulty can arise when the phenomena of interest is itself involved in a modelling or learning process of some kind. These include what Rosen called “anticipatory systems” [3]. These include organisms, of course, but also simpler systems where there is some adaptation process occurring. The fact that the system is adapting, learning or modelling is not necessarily a problem, but if the feedback involved becomes very complicated (in the sense that future adaptations are somewhat dependent upon past adaptations etc.) then this can act to create a new context in itself. That is, what is a valid model is more dependent upon the intricacies of the adaptation process than factors external to the adaptation process so the cluster of model scopes is internal to that process, rather than something with so much reference to its environment (though it may be a combination of both).

This is a relative difficulty. For fairly simple adaptive processes the adaptive process may be fairly predictable given other conditions and so be modellable from the context that the process exists in. But adaptation, under some circumstances has the effect of embedding causation endogenously, and this may prevent its clean separation from an external viewpoint. Evolution and evolution-like processes seem to result in complex adaptations and adaptive mechanisms, with the result that the internal processes within organisms (and cells) effectively form their own contexts (as observed in [9]). Similarly the embedding through evolutionary adaptation of species into ecosystems has a similar effect making the niche or the ecosystem the effective context. In these cases it is very difficult to model across these contextual boundaries due to the “density” and “criticality” of the interactions within these contexts. It may

be, that in many cases, we have to be satisfied with multiple but separate models for the cellular and ecological contexts, with only informal relationships between them.

Conclusion

I hope the sketchy account above, can convince readers that there is nothing to be afraid of in the thought of Robert Rosen. I think his arguments were slightly distorted by his frustration at not being understood, and his understandable reaction in trying to *prove* he was right and escape the scientific rigidities of his day. However, as recast in the light of the context-dependency of causation and hence the context-dependency of causation formal modelling of complex phenomena, it can be used as the start of a development of a more rigorous and careful scientific methodology, that uses multiple models and different contexts in a complementary way, whilst being aware of some of the conditions where the models are liable to fail.

References

1. R. Rosen, *Life Itself - A Comprehensive Enquiry into the Nature, Origin and Fabrication of Life*. **1991**, New York: Columbia University Press.
2. R. Rosen, R.. Organisms as Causal Systems Which Are Not Mechanisms: An Essay into the Nature of Complexity. In *Theoretical Biology and Complexity: Three Essays on the Natural Philosophy of Complex Systems*. London, Academic Press, **1985**, p. 165-203.
3. R. Rosen, *Anticipatory Systems*. **1985**, New York: Pergamon.
4. R. Rosen, Bionics Revisited, in *The Machine as Metaphor and Tool*. **1993**, Springer-Verlag: Berlin. p. 87-100.
5. B. Edmonds, *Pragmatic Holism*. *Foundations of Science*, **1999**, 4, 57-82.
6. M. B. Hesse, *Models and Analogies in Science*. **1963**, London: Sheed and Ward.
7. P. Hayes, Contexts in Context. In *Context in Knowledge Representation and Natural Language*, AAAI Fall Symposium, **1997**. AAAI Press.
8. B. Edmonds, The Pragmatic Roots of Context, In *Proceedings of the Second International and Interdisciplinary Conference on Modeling and Using Context*. **1999**, 1688, 119-134.
9. W. Zadrozny, *A Pragmatic Approach to Context*. In *Context in Knowledge Representation and Natural Language*, AAAI Fall Symposium, **1997**, MIT. AAAI Press.
10. B Kokinov and M. Grinberg, *Simulating Context Effects in Problem Solving with AMBR*. In Akman, V., et al. (eds.) *Modelling and Using Context*, Springer-Verlag, *Lecture Notes in Artificial Intelligence*, **2001**, 2116, 221-234.
11. J. Pearl, *Causality - Models, Reasoning, and Inference*, Cambridge University Press, **2000**.
12. J. A. Fodor, Special sciences (Or: The disunity of science as a working hypothesis). *Synthese*, **1974**, 28, 97-115.
13. D. Rabounski, Zelmanov's Anthropic Principle and Infinite Relativity Principle, *Progress in Physics*, **2006**, 1, 35-37.

14. B. Edmonds, *The Nature of Noise*, Epistemological Perspectives on Simulation - II Edition, University of Brescia, Italy, October 2006. <http://cfpm.org/cpmrep156.html>.
15. K. Kaneko, *Globally Coupled chaos violates the law of large numbers but not the central limit theorem*. *Physical Review Letters*, **1990**, 65, p. 1391-4.
16. B. Edmonds, *Modelling Bounded Rationality In Agent-Based Simulations using the Evolution of Mental Models*, in *Computational Techniques for Modelling Learning in Economics*, T. Brenner, Editor. **1999**, Kluwer. p. 305-332.
17. A. M. Turing, *On Computable Numbers, with an application to the Entscheidungsproblem*, *Proc. London Math. Soc.*, **1936**, 42, 230-65.
18. B. Edmonds, *The Constructability of Artificial Intelligence (as defined by the Turing Test)*, *Journal of Logic Language and Information*, **2000**, 9, 419-424.
19. R. N. Giere, *Explaining science: a cognitive approach*. Chicago; London, University of Chicago Press, **1988**.
20. N. Cartwright, *How the Laws of Physics Lie*. **1983**, Oxford: Clarendon Pres.
21. B. Edmonds, *Capturing social embeddedness: A constructivist approach*. *Adaptive Behavior*, 1999, 7(3-4), 323-347.