



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Altered States and Virtual Beliefs

Citation for published version:

Oberlander, J & Dayan, P 1996, Altered States and Virtual Beliefs. in A Clark & P Millican (eds), *Connectionism, Concepts and Folk Psychology: The Legacy of Alan Turing*. vol. 2, Mind Association Occasional Series, Oxford University Press, USA, pp. 101-114.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Connectionism, Concepts and Folk Psychology

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Altered States and Virtual Beliefs

Jon Oberlander*

Peter Dayan†

April 1990

Abstract

Functionalism avoids a potentially fatal infinite regress by realising the low level phenomena of mind in mere Turing machines rather than via undischarged homunculæ. Adopting a narrow, logical view of such mechanisms has led to two sorts of problems: exponentially intractable models of human inferential competence; and a wide divergence between the supposed competence and observable human performance. These problems have produced a retreat to instrumentalism about folk psychological concepts—a retreat that has long been trumpeted by certain philosophers.

We argue that abandoning hope in this way is premature. The problems arise from two quarters: an impoverished notion of competence models that divorce them too far from performance models; and too narrow a view of possible functional mechanisms. We explore the consequences of retaining a traditional view of inference, while adopting a new mechanism for memory. We motivate the latter by developing a seemingly unnatural picture of von Neumann machines, taking their finite memory seriously. Together, the two repairs suffice to make instrumentalism avoidable.

1 Introduction

Functionalism derives its explanatory force from its firm grounding in the certainties of Turing machines (TMs) and their practical approximations, the von Neumann machines (vNMs). One branch of cognitive psychology, and its philosophical stablemate, has taken this doctrine seriously, and coupled it with some of the basic notions of folk psychology such as beliefs, desires, hopes, *etc.* Most of the theories are wedded rather firmly to these architectures, with all their connotations of logicity. They therefore face problems explaining the apparent *illogicality* and obvious *tractability* of human cognition. In desperation, some functionalists, whilst naturally retaining the Turing metaphor at the heart of their theories, have been forced to abandon these folk psychological constructs. The switch is either to instrumentalism, proposing that these constructs are *not* causal in producing behaviour but are still useful for explaining or predicting it [4], or else to a more radical position, holding that they are not even satisfactory for that [23].

*Human Communication Research Centre, University of Edinburgh. This work was carried out with the support of the British Academy. Email address: jon@cogsci.ed.ac.uk@nsfnet-relay.ac.uk

†Centre for Cognitive Science and Department of Physics, University of Edinburgh. Supported by the SERC.

Connectionism, for all its many faults, gives the first hint that something may be wrong with this retreat. It demonstrates that, despite essential computational equivalence, the strict analogy with TMs that has survived into vNMs is not the only way of conceiving of the functional mechanisms of thought. Furthermore, on broadening the range of possible mechanisms, it becomes apparent that too simplistic a notion of human competence is potentially dangerous. Unfortunately, connectionism seems rather remote architecturally and representationally from high level psychological entities, so some *independent* motivation behind alternative functional mechanisms is required.

Section 2 works through some of the reasoning behind instrumentalism, involving such problems as virtual beliefs and inferential tractability, and discusses certain conditions under which the instrumentalist retreat may be avoidable. Section 3 sketches out a view of functionalist mechanisms that satisfies these conditions. Section 4 concludes that the case for abandoning folk psychology, an unattractive move in the absence of a viable alternative, is not yet proven.

2 Competence, Performance and Virtual Belief

The process of ‘systematising’ folk psychology involves constructing a functionalist tableau around its basic entities such as beliefs and desires, and its basic notions of how they interact. In cartoon form, one such notion might be ‘if person A believes that doing B will result in C, and desires that C, then, all else being equal, A will do B’. Folk psychology is not systematic in any functionalist sense, and does not concern itself with mechanism or the causal efficacy or otherwise of its entities.

The overall tableaux most commonly suggested involve sentential belief boxes and desire boxes containing the beliefs and desires respectively, together with central processors that execute the ‘rules’. Although this is but a crude sketch, it has two key features that are common to most such postulations: the one-by-one consideration of the folk psychological entities; and the variously expressed commitment to some variety of logic for the rule processing. As hinted above, both of these are strongly associated with the traditional view of the vNM.

It is important to remember that one cannot attack **all** systematisations by merely attacking one. In fact, the ongoing debate about folk psychology between Stich [23], Dennett [5], Fodor [9], and many others, is conducted *solely* about the kind of model just outlined. Opponents of the *Language of Thought* (LoT) type doctrine point to various flaws and apparent absurdities in the inferences it suggests. They then either reject folk psychology totally, or else suggest that its efficacy is merely approximate, and that its entities are not real, extant and causal, but merely abstract calculating devices with which one can predict behaviour. The former view seems a little churlish given that folk psychology seems on the face of it to be the only theory that has any predictive power. The latter view is at root unattractive since it is not accompanied by any compelling alternative causal entities. However, if the problems for folk psychology suggested by Dennett and others are indeed damaging in the way posited, instrumentalism would be the only course open. We shall proceed by establishing and examining the framework within which the objections are cast.

2.1 Competence and Incompetence

The first objection comes from a concern about the mechanistic relevance of Chomskian [2] competence theory rather than a performance theory. On the one hand, there is the competence of the undistracted, resource-unlimited ideal thinker. On the other, there is the performance of distracted, error-prone real thinkers. Dennett (1987:76) has emphasised the significance of the correct choice of competence model for psychological theorising. If we get the competence model badly wrong, any performance model based on it will be full of artefactual results and *ad hoc* mechanisms. Poor results with this performance model will tend to reflect back on the competence model. Dennett says that:

The fact about competence models that provokes my “instrumentalism” is that the decomposition of one’s competence model into parts, phases, states, steps or whatever *need* shed no light at all on the decomposition of actual mechanical parts, phases, states, or steps of the system being modeled—even when the competence model is, as a competence model, excellent. [Dennett 1987:76]

Any competence model that suggests that human reasoning is fundamentally logical runs up against problems such as those from human conditional reasoning. In some tasks [20, 14], 97% of humans fail to satisfy the logical idealisation, so the LoT might in fact rather be called an *In*competence model. This is not in itself destructive of the LoT, since it is always possible to multiply entities to save the phenomena. Further arguments below about virtual belief are more persuasive.

It is important to realise that the truth of the quotation and the unpalatable facts about logical reasoning are not arguments against **all** possible systematised folk psychology, only against the LoT version. Of course, competence theories may say nothing about performance mechanisms, but that does not absolve them from empirical and theoretical study. If, as in the case of purely logical inference, incompetence is found to be rife, and furthermore achieving competence is known to be computationally intractable, it follows that the competence model is incorrect. Note also a possible lacuna made apparent by this; if the competence model is too far divorced from performance, it is useless even for instrumental theorising, since its predictions will be invalid.

2.2 Virtual Belief

Dennett does not reject the LoT view on account of conditional reasoning data. He prefers rather to look at the more complex case of virtual beliefs. They provide an interesting test case in view of the alternative explanation available from the memory machine picture that will be established in Section 3.

When describing what someone believes, we can distinguish at least three significantly different cases. Take Neil. He believes that Margaret is Prime Minister. No doubt he also believes that $30001 > 30000$. And we judge, watching him play chess, that he believes that you should get the queen out early. In making these ascriptions, we have traded on three different kinds of belief. The belief about Margaret is probably something that Neil has actively considered and has stored (after the 1987 general election, perhaps). As such, it is perhaps an *explicit* belief. The belief about the numbers he probably never entertained

before, but would assent to immediately if asked. It's something that his other beliefs, about numbers and arithmetic, commit him to. Indeed, it could perhaps be inferred, via a laborious logical proof, from what he explicitly believes. Let's call this an *implicit* belief. Given this ostensive definition, it should be obvious that Neil is committed to an infinite number of such implicit beliefs. Finally, his belief about chess is something he has never entertained before, that he might not admit, and perhaps cannot infer from his explicit beliefs. It too is implicit in the way he behaves. But to distinguish it from the second case, let us call it an *emergent* belief. Implicit and emergent beliefs are sometimes known as *virtual* beliefs.

Now, if we assume the LoT model of inferential behaviour, we must constrain our theorising about possible performance models to be compatible with both the competence model, and with some of the hard facts about explicit and virtual belief. Since reasoning is fundamentally logical, but humans rarely seem to believe all the consequences of their beliefs, it is tempting to build a performance model of reasoning along the following lines.

Explicit beliefs form a core library of axioms. A logical inference device, commonly called the Extrapolator-Deducer (XD), derives implicit beliefs from these axioms by applying the internal rules of inference [9, 4]. If you ask Neil about the number 30001, he will be able to derive the belief in question very rapidly. Emergent beliefs don't quite fit into this picture, and proposals have been made that they are not really represented, either statically (as explicit beliefs) or on demand (as implicit beliefs), but are really "there" as properties of the architecture of the belief retrieval system *eg* [16].

This is a performance model made in hell. Various problems with it are well-known. The XD must have its own stock of beliefs in order to know which beliefs to use in deduction. The power of the XD will vary from person to person. The number of inference steps the XD can take in (say) 100ms is arbitrarily fixed at (say) 15—but we can't tell which logical system the XD is using. There are some implicit beliefs that the XD should be able to infer that it can't, and *vice versa*. Further, it is really not at all obvious that an account of emergent beliefs can be grafted onto the XD. In what sense does the "Get the queen out early" belief reside in the activity of the XD?

Dennett expresses these problems with performance models of virtual belief effectively in his description of a machine designed to get jokes about Newfoundlanders. He supplies a joke (but attributes it to Pylyshyn), and lists various propositions we have to believe if we're to get the joke, and states:

Not only do these beliefs not "come to mind" . . . , but it is also highly implausible to suppose that each and every one of them is consulted, independently, by a computational mechanism designed to knit up the lacunæ in the story by a deductive generation process. . . . The list of beliefs gives us a good general idea of the information that must be in the head, but if we view it as a list of axioms from which a derivational procedure deduces the "point of the joke," we may have the sketch of a performance model, but it is a particularly ill-favored performance model. [Dennett 1987:77]

Following this argument leaves only two options to those who insist on rejecting all systematised folk psychology on the basis of the performance inadequacies of the LoT; one can either retain folk psychology (and *a fortiori* its entities such as belief) as an instrumental,

non-causal, theory, or one can reject it altogether. A third option, which is followed below, is to suggest an alternative systematisation.

It bears notice that Dennett's preferred option, the first of the above, involves a certain prestidigitation. If folk psychology is to be useful to us in predicting and explaining other people's behaviour, we must have a way of using *their* assertions of beliefs and desires appropriately. If the theory is so incoherent, intractable, or implausible, that these entities cannot be implicated in **causing** the behaviour, how can they be implicated in **explaining** or **predicting** it? Dennett also faces criticism from the radical anti-folk psychologists such as Stich [23], who argue that having rejected causality, to retain beliefs and desires *etc* in explaining behaviour is like explaining planetary motion in terms of Ptolemæan epicycles whilst rejecting their entire basis.

Ruling systematised folk psychology out of court is a fundamentally unattractive course, since there is no equally predictive, let alone better, replacement. Such a move can be avoided if some alternative XD theory can be developed, particularly one that has independent motivation. This is the aim of the next section.

3 Altered States

In searching for new and fruitful analogies with which to extend and defend their models of cognition, functionalists continually find their writhing room limited by fundamental Turing equivalence. Despite the possibilities of re-implementation and simulation of one such computational system by another, properties that have natural explanations in the one may seem mysterious in the other. By applying, and then extending, one such transformation, we will attempt to demonstrate this constructively, and to produce a model of memory that can be used in an alternate extrapolator-deducer theory.

The obvious evolutionary path from the TM architecture to the vNM architecture is to leave unchanged the nature and rôle of the central processor (the reader) and its states, but to regard the memory of the vNM as a *finite* tape. Arbitrary locations on this virtual tape can be accessed in $O(1)$ rather than $O(N)$ time, where N is its virtual length. The ability easily to write to, or read from, arbitrary locations makes a vNM considerably easier to program than a TM, and the (almost) size-independent read/write times, makes a large class of algorithms computationally feasible. However, the recent debate pitting Reduced Instruction Set (RISC) against Complex Instruction Set (CISC) microprocessors points to an obvious trade off that, as described, is rather hidden in the detail of Turing universality. Specifically, there is a balance between the complexity of the set of states and state transitions of the central processor on the one hand, and the complexity and length of the programs that perform a task, on the other.

Since the memory of a vNM is finite, one can take this view on complexity further. We can draw the 'state *versus* tape' boundary around not only the central processor, but also some or all of the memory. Under the new description, states represent not only the contents of the internal flags and registers of the microprocessor, but also the entire contents of memory—roughly, all the transistors that comprise the operation of the machine. This new description is precisely that of the entire computer as a finite state machine (FSM), which of course is exactly what its finite memory permits it to be. State transitions are instigated just

as before by the operation of the central processor, but *now* the state transition diagram is extremely rich and complex. The way that von Neumann machines are designed has the effect of enforcing what might be described as *locality* in this state transition diagram. That is, if the states were ‘ordered’ in an appropriate way,¹ most of the possible next states themselves would be local.

The new complexity and richness of the set of states and state transition diagram in this alternative viewpoint accounts for the infrequency of its adoption. But it is precisely the *attraction* for us. The distinctions between program and storage appear considerably eroded, and we will go on to erode them even further. Note that at this stage this is only an alternative perspective of exactly the same machine, which is condemned to traverse exactly the same states and execute in the same way exactly the same programs.

Various key properties of computation look very different on this alternative picture of the vNM, particularly symbolic representation and the potentially holistic influence of memory. What could previously be identified as separate contents of memory can only be found, if at all, in the **labels** of a whole set of states; and the labels, of course, have no causal rôle in engendering the behaviour of the system. It is therefore no longer even sensible to try to isolate out particular symbol structures stored at particular places in memory, or to determine the nature and provenance of their logical effects. The symbol processing properties of the whole machine are entirely emergent from the dynamics of the simple low level entities.

In addition, the whole of memory is wrapped up in a state, so moving from one state to another could correspond to a radical change in its entire contents. Of course, the vNM architecture prevents such transitions from happening. Memory is like tape, inviolate except serially. A rational explanation for this severe constraint is that such radical changes are potentially dangerous in the absence of any well-founded method of policing them. However, different mechanisms for state transition need not necessarily labour under such constraints.

Removing another such bar could allow the whole contents of memory not only to be changed by such a state transition, but also to determine which state transition occurs. This new picture naturally encompasses a much richer theory of the interaction between memory and inference, mediated in a complex fashion by the hardware of the machine. It would also appear to be more commensurable with the views on cognition of Maturana and Varela [18]. They consider the nervous system less as a machine responding as a program to a set of inputs with a set of outputs, and more as a system whose states and dynamical state transitions are affected by all its components. Some of these will have characteristics partly determined by events in the external world.

We do not have, and indeed to Dennett do not need, a sophisticated picture of how to develop from the above picture a complete XD. The naive mechanism which will suffice may best be imagined as some form of **PROLOG** system whose inferences are executed on one computer, but whose database and short term memory are stored on another. The first requests rules and facts from the second, and demands that changes be made in the light of its inferences. Although continuing to view the first computer, the inference engine, in a traditional way, though not of course as a **PROLOG** machine, we will investigate the

¹By the ‘invisible’ values of the program counter, the registers and flags, the possible values of the memory location pointed to by that program counter and then all the other possible values of all the other memory locations.

consequences of considering alternative mechanisms underlying the function of the second computer, the memory machine, by further developing the above picture. This general view of the relationship between memory and inference in an XD has a number of defenders, *eg* Anderson [1].

Consider now the memory machine (MM). Under a traditional description it would contain symbols, and would be searched (subject to a request from the inference engine) according to some criteria, such as ‘find a script containing doctors’, or ‘what is the capital of France?’, or the like. Under the new description, the answers to such queries can be determined from the state the MM occupies, and its own searching procedure for finding the required information can be seen in the set of states traversed. However, on this picture, it is natural to contemplate altering the hardware of the machine, the complex system that applies the rules for moving from one state to another, and consequently changing the whole *modus operandi* and output of the MM.

One such change is to consider the states as the local minima of some function (perhaps an *energy function*, for want of a better term), and the state transitions as movements from one local minimum to another (*cf* [13, 10]). The hardware can be considered as moving the system around the energy surface, according to its principles, landing up in states at the local minima. Learning changes the energy surface so as to alter either the minima, or the traversable paths from one minimum to another, and consequently the states actually entered (\sim the memories recalled) by the MM under given conditions. This, as advertised, licenses state transitions *other* than those allowed by the vNM. Now, the whole contents of memory are causally implicated in every move.

The appeal of making this change is that it’s much easier to see how to build a human-like memory (if not actually to do so). Contextual flexibility, associational memory, and inferential error, can all be given natural explanations on this model. Of course, we are not able to specify either the energy function, or the mechanisms that are sensitive to it, but in that respect we are no worse off than other functionalists, and need only demonstrate possible Fodorian ‘floating straws’.

The evolutionary sequence has been as follows. First we had the Turing machine. Then we had the von Neumann machine as the implementation of a finite TM. Then we had the vNM as a finite state machine. And finally we reached alternative mechanisms for state transition in a FSM, and consequently alternative state transitions. This motivates and locates a (connectionist) dynamical picture of the mechanisms of mind. It also neatly demonstrates the difficulties connectionists are bound to have in talking about symbolic representation in their models. Even when we know that the symbols are explicitly emergent, as in the FSM description of the vNM, they are elusive, to say the least.

We have previously [3] linked related notions of local **maximisation** (reversing the sign) with those of Elster [7], Doyle [6], and Simon [21]. Very briefly, the suggestion was that global maximisation based on predictions of the future, which Elster claims to be criterial for rationality, is too difficult for mere machines (including us) in interesting domains. However, not all local maxima are as high as others, and one can understand Simon’s satisficing, and consequently Doyle’s ideas of rational psychology as utility maximisation, in these terms. These maxima are actually on a longer and larger scale than the ones on which we concentrate in this paper, since here, the local maxima just define the raw material for the inference engine. However, large scale **global** maximisation cannot be based on small

scale **local** maximisation, which is the basis for our claim about the difficulty.

It is worth very briefly comparing the possible alternative views on representation. Symbols can vary on two dimensions: hard *versus* soft, and active *versus* passive. Hard symbols are like those to be found on a tape of a TM, in the memory of a vNM, or indeed in the localist models of Hofstadter [11, 12] or Feldman and Ballard [8]. They are hard in the sense that their semantic (and usually syntactic) boundaries are fixed at ‘compile’ time and do not vary during the course of a computation. Soft symbols are more like those of distributed connectionist models, or (hopefully) the emergent symbols of the MM described above. Their mutual semantic boundaries vary in a context-dependent way, and indeed are usually learnt during the operation of the system. The distinction between passive and active symbols is rather less precise, since it gets unhelpfully entangled with the differences between declarative and procedural representations. Passive symbols are like those in a typical frame-type system [19], in that they are more ‘processed over’ than ‘processing’ entities. Active symbols are just the reverse, as in Hofstadter’s fluid concepts or the holistic MM above. Adequate context-dependency and learning seem to require soft, active symbols, since all attempts at passive symbols seem to falter on the former criterion, and hard ones on the latter.

This section has attempted to motivate a significant move away from the traditional model of computation by taking seriously the obvious fact that a von Neumann machine has only finite memory. The reconceptualisation involved can lead to an architecture where states are actually large dimensional vectors upon which soft symbols are emergent, and whose state transitions correspond to movements along some form of energy surface defined by the interactions between the various components in the system. The hard symbols of Turing machine tapes and the active symbols of Hofstadter’s and Feldman and Ballard’s models are both rejected.

We are now in a position to evaluate the consequences for the behaviour of the overall XD, like Dennett’s joke-getting system. We have changed the functioning of the MM, which of course stores both the rules consulted by the IM, and the memories used in executing the rules. The contextual flexibility built into the MM derives from the implication of the entire contents of memory in state transitions. Even in the naive picture which has only those entities “come to mind” that are passed from the MM to the IM, it is apparent that the implicit Newfoundlander-joke beliefs that gave Dennett so much trouble need never be passed into the IM. Instead, they can be causally implicated in the overall XD process by helping determine which consciously ‘observable’ beliefs are selected during the non von Neumann search.

Implicit beliefs generally gain a more natural explanation on this picture, since an intractable search amongst rules or memories is not necessary. Indeed, explanations can be imagined for more complex cocktails of repressed beliefs and their apparent effects in terms of modifications to the energy surface determining the search. Of course, this story is incomplete in the absence of an exact specification of the requisite ‘all-singing, all-dancing’ state-transition system, but work on connectionist memory systems demonstrates at least its potential feasibility, and that is all that is required.

4 Conclusions

In endorsing functionalism, it has traditionally been too easy to adopt a narrow view of computational mechanisms. The fact that two architectures may be equivalent in the Turing sense does *not* imply that the functional properties that can naturally be ascribed to one can naturally be ascribed to the other. Attempts to systematise folk psychology based on the narrow, logical, view have failed. They have failed because of the obvious illogicality and tractability of human reasoning, and because of the implausibility of the inferences which the narrow view insists on. Instrumentalists with the *same* narrow view of mechanism conclude from this failure that the entities and notions of folk psychology are inherently wrong, and that all they can do is provide an abstract, non-realistic, tool for computation.

An alternative conclusion is that it is not the entities and notions of folk psychology that are at fault. It is their *particular* systematisation under the narrow view. A broader perspective can be adopted, one that can be motivated directly by two moves. First, we draw an alternative boundary between the central processing unit and the memory of a von Neumann machine regarded in finite state terms. Secondly, we modify slightly the state transition rules. Instrumentalism was at best unattractive because of its non-causal entities. With richer functional mechanism to hand, we can now reject it as an *unnecessary* evil.

References

- [1] Anderson, J. R. (1983) *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.
- [2] Chomsky, N. (1965) *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- [3] Dayan, P. & Oberlander, J. (1989) *Maximal Mechanical Rationality*. Centre for Cognitive Science, Edinburgh.
- [4] Dennett, D.C. (1978) *Brainstorms: Philosophic Essays on Mind and Psychology*. MIT Press/Bradford Books, Cambridge, MA.
- [5] Dennett, D.C. (1987) *The Intentional Stance*. Cambridge, MA: MIT Press/Bradford Books.
- [6] Doyle, J. (1983) *Some theories of reasoned assumptions: an essay in rational psychology*. (Tech. Rep. No. CMU-CS-83-125). Pittsburgh, PA: Carnegie-Mellon University. Department of Computer Science.
- [7] Elster, J. (1984) *Ulysses and the Sirens: Studies in rationality and irrationality*, Revised Edition. Cambridge: Cambridge University Press.
- [8] Feldman, J.A. & Ballard, D.H. (1982) Connectionist models and their properties. *Cognitive Science* **6**, 205-254.
- [9] Fodor, J.A. (1975) *The Language of Thought*. Hassocks, Sussex: Harvester Press.
- [10] Hinton, G., Sejnowski, T. & Ackley, D. (1984) *Boltzmann machines: Constraint satisfaction networks that learn*. (Tech. Rep. No. CMU-CS-84-119). Pittsburgh, PA: Carnegie-Mellon University. Department of Computer Science.

- [11] Hofstadter, D.R. (1979) *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, NY.
- [12] Hofstadter, D.R., Mitchell, M. & French, R.M. (1987) *Fluid Concepts and Creative Analogies: A Theory and its Computational Implementation*. FARG DOC 87-1. University of Michigan. Ann Arbor, Michigan.
- [13] Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, USA*, **79**, 2554-2558.
- [14] Johnson-Laird, P. N. & Wason, P. C. (1972) *Psychology of Reasoning: Structure and Content*. Cambridge, MA: Harvard University Press.
- [15] Lycan, W. G. (1981) Form, function, and feel. *Journal of Philosophy*. **78**, 24-50.
- [16] Lycan, W. G. (1988) *Judgement and Justification*. Cambridge: CUP.
- [17] Marr, D. (1982) *Vision*. Freeman, NY.
- [18] Maturana, H. & Varela, F. (1972) *De Maquinas y Seres Vivos*. Editorial Universitaria, Chile. Published in English (1980) *Autopoiesis and Cognition: The Realization of the Living*. Reidel, Dordrecht.
- [19] Minsky, M. (1975) A framework for representing knowledge. In Winston, P. ed. *The Psychology of Computer Vision*. Englewood Cliffs, NJ: Prentice-Hall.
- [20] Oaksford, M. R. (1988) *Cognition and Inquiry: An Investigation into the Psychosemantics of Conditional Reasoning*. PhD Thesis. Centre for Cognitive Science, Edinburgh.
- [21] Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*. **69**, 99-118.
- [22] Smolensky, P. (1988) On the proper treatment of connectionism. *Behavioural and Brain Sciences* **11**, 1-74.
- [23] Stich, S. (1983) *From Folk Psychology to Cognitive Science: The Case Against Belief*. MIT Press, Cambridge, MA.
- [24] Winograd, T. & Flores, F. (1986) *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publisher, Norwood NJ.