The Limits of Mechanistic Explanation in Neurocognitive Sciences

Anna-Mari Rusanen (anna-mari.rusanen@helsinki.fi)
Department of Philosophy, PO Box 9
00014 University of Helsinki, Finland

Otto Lappi (otto.lappi@helsinki.fi)
Cognitive Science Unit, Department of Psychology
PO Box 9, 00014 University of Helsinki, Finland

Abstract

In this article, we discuss some aspects concerning the appropriate model of explanation for the emerging neurocognitive sciences. The question we pose is whether computational explanation is a form of mechanistic explanation, or whether there are characteristics that go beyond, in important ways, the explanatory framework of mechanistic explanation. We will argue that if one considers neurocognitive systems at a genuinely cognitive level, where explanations refer to the level of description that Marr called computational, the mechanistic approach, as it stands, does not offer an appropriate model of explanation.

Introduction

It has been suggested by many philosophers that the mechanistic model of explanation could perhaps offer a satisfactory model of explanation for the neurosciences (Bechtel and Richardson, 1993; Bechtel and Abrahamsen, 2005; Craver, 2005; Craver and Darden, 2005; Glennan, 2005; Machamer, Darden and Craver, 2000). Quite recently it has been also proposed that the mechanistic model could be extended to cover computational explanations as well (Piccinini, 2004; Piccinini, 2006a; Piccinini, 2006b).

In this article we explore some aspects of this mechanistic model of explanation in the context of neurocognitive sciences. By "neurocognitive sciences" we mean those branches of theoretical and computational neuroscience that are explicitly concerned with cognitive, i.e. information processing, capacities of organisms. These include, for example, traditional cognitive neuroscience which is based on neuroimaging and theories adopted from cognitive psychology, as well as those strands of computational neuroscience that attempt to model brain functions at the cognitive level (for the distinction see Grush, 2001). We are therefore not concerned here with those strands of computational and systems neuroscience that are not in the business of computational explanation of cognition, but are instead trying to discover the principles of organization of the brain operating at levels of explanation "lower" than that required to account for cognitive organization. These principles of organization may be small scale –e,g, molecular or cellular–or large scale-i.e. systems level or behavioral. The crucial issue is, however, whether explanations make reference to computations performed by ot implemented in the brain.

We agree with "the mechanists" that the mechanistic model of explanation is useful in neurosciences on levels of description higher than physiology, at least up to the "functional" level of molecular and systems neuroscience. We will even be prepared to assume with the mechanicists that up to the "syntactic" or "algorithmic" levels, it is possible for the purposes of computational explanation in neuroscience to describe and explain computational processes mechanistically (cf. Piccinini, 2004; Piccinini 2006a; Piccinini, 2006b). Thus we agree that mechanistic explanations are an important part of the explanatory framework of the neurocognitive sciences.

However, we still maintain that if one considers the neurocognitive systems at a genuinely cognitive level, and specifically consider neurocognitive explanations which refer to the level of explanation that Marr (1982) called computational, the mechanistic approach, as it stands, does not offer an appropriate model of computational explanation. (This form of explanation is, furthermore, characteristic of the cognitive sciences. It has no direct analogy in the non-cognitive non-representational branches of bio-, computer- or neurosciences for which the mechanistic model of explanation was mainly developed). Thus, it may be that the mechanistic model of explanation cannot capture that which is special in explanations in cognitive, and neurocognitive sciences generally.

What we are concerned, then, is to what extent this approach to "computational mechanisms" can be subsumed within the framework of mechanistic explanation.

The Mechanistic Model of Explanation

We shall concern ourselves with the question of whether or not models of neurocognitive explanation in computational theories of the neurocognitive sciences can be satisfactorily described within the mechanistic framework of explanation. But first, we need to make more precise what we mean by a mechanism and a mechanistic explanation.

The most recent discussion of mechanistic explanation has emerged in the context of philosophy of the life sciences (Bechtel & Abrahamsen, 2005; Bechtel & Richardson, 1993; Craver, 2005; Craver & Darden, 2005; Glennan, 2005; Machamer, Darden & Craver, 2000). Explanations in those sciences usually involve presenting a model of a mechanism that is taken to be responsible for a given phenomenon (Bechtel & Richardson, 1993; Craver, 2005;

Glennan, 2005). This is meant to be in contrast to deriving the phenomenon from universal laws and initial and boundary conditions.

In Craver`s (2001, 2005, 2006) model a mechanism is a structure performing a function in virtue of its component parts, component operations performed by those parts, and the organization of the parts into a functional whole (the "system"). For example, the heart's function is to maintain blood pressure and circulation. This is achieved by an intricate system contractile fibers and a neural mechanism to synchronize their contractions. (Craver distinguishes etiological, contextual and constitutive mechanistic explanations (2001, 2005). In this article we focus only on constitutive explanations).

Constitutive mechanistic explanations are explanations in which the explanandum is always at higher level of explanation than the explanans. The explanandum is the activity or behavior exhibited by a system, and the explanans is a description of the internal organization or structure of the system, in terms of the lower level entities, activities and their relations which together constitute the system and are causally responsible for its functioning. The explanandum of a constitutive mechanistic explanation is some behavior of the mechanism or some output caused by or organized by the mechanism (Craver, 2005; Craver, 2006; Machamer, Darden & Craver, 2000). In the case of the heart, for example, the pumping action and ensuing circulation of the blood are the phenomena to be explained. The explanans of a phenomenon is a model of the mechanism that describes the causal agents responsible for carrying out the component operations that produce the phenomenon (Craver, 2006).

From the mechanistic perspective, computational mechanisms are a special case. They are mechanisms that can be made to perform information processing (rational and semantics-sensitive) tasks. In other words, computational mechanisms are devices where the design allows one to set initial conditions so that when some piece of information from a domain is encoded in them, they will operate on that information and, generally, when the mechanism is done, one ends up with another piece of information from the same domain, in some interesting way semantically related to the first piece.

Recently Piccinini (2006a, 2006b, 2006c) has proposed that computational explanation is mechanistic, since the computing mechanisms - such as the microprocessors of desktop computers or cell assemblies communicating via action potential spikes - are analyzed in terms of their component parts, their functions, and their organization Computational explanation is then "a mechanistic explanation that characterizes the inputs, outputs, and sometimes internal states of a mechanism as strings of symbols, and it provides a rule, defined over the inputs (and possibly the internal states), for generating the outputs" (Piccinini, 2006b).

Piccinini holds the mechanistic view of computational individuation, i.e. that computational states are individuated by their mechanistic properties without appealing to any "higher level" semantic properties. These properties are supposed to be picked out by the appropriate mechanistic explanation (Piccinini, 2006a). In order to draw the distinction between relevant and irrelevant properties, one needs to know "which of a computing mechanism's properties are relevant to its computational inputs and outputs and how they are relevant" (Piccinini, 2006a). Piccinini suggests that the relevance (for computing) of a certain mechanism would be evaluated by knowing how the mechanism`s inputs and outputs interact with their context (Piccinini, 2006a), i.e. by knowing its causal (as opposed to, say, intentional) relations with the environment.

We disagree. In neurocognitive sciences, a computational mechanism is evaluated *for explanatory relevance* at least partially by describing the *task* of the computing mechanism. The theories, which specify the nature of neurocognitive tasks or functions are theories of neurocognitive competences, or "computational level theories" in Marr's sense of computation (as opposed to representation and algorithmic mechanisms). In this approach, the computational level is invoked to explain and assess the explanatory relevance (but not necessarily to individuate) mechanisms at the algorithmic and implementation levels. Mechanism individuation cannot be done in the sense Piccinini envisions.

But even if it could, there is another problem for extending mechanistic explanation to computational explanation, that of the direction of interlevel explanation.

In the mechanistic approach, interlevel explanation is "bottom-up", in that the (relevant) mechanistic organization explains the competence. In the next section we will take a look at the marrian approach to computational theories of neurocognitive competences, where explanations can also be "top-down". We take it that the track record of this sort of approach (in the domains of vision, language, and the probabilistic approach to cognition) is at least strong enough to merit philosophical attention. There is serious and interesting work being done in theoretical neuroscience and cognitive modeling within this framework (for overviews, see Anderson, 1991a; Anderson1991b; Chater, 1996; Chater et al, 2006; Friston 2002; Friston 2005).

Marr and computational explanation

In Marr`s taxonomy the computational level is the level that gives an account of the tasks that the neurocognitive system performs, or problems that the cognitive system in question is though to have the capacity to solve, as well as the information requirements of the tasks (1982). This level is also the level, whereby the appropriateness and adequacy (for the task) of mappings from representations to others are assessed (cf. Marr, 1982).

In the case of human vision, one such task might be to faithfully construct 3D descriptions of the environment from two 2D projections. The task is specified by giving the abstract set of rules that tells us what the system does and when it performs a computation. This abstract computational theory characterizes the tasks as mappings, functions from one kind of information to another. It constitutes, in other words, a theory of competence for a specific cognitive capacity - vision, language, decision making etc.

Descriptions of performances corresponding to particular competences are given at the algorithmic or implementation levels; in the case of neurocognitive sciences these might go all the way down to functional or neurological (systemic, cellular or molecular) level models. In a nutshell, cognitive *competences* are specified at the higher computational level in Marr`s taxonomy and *performances* at the lower, algorithmic and implementation levels.

It is important to acknowledge that at the computational or competence level theories one does not yet commit to anything specific concerning the mechanisms carrying out the required information processing. The specification of a competence is abstract in the sense that nothing specific in terms of entities or activities causally related to each other is yet assumed.

The computational theory answers "what"- questions and "why" questions. These questions are typically formulated as questions like "what is the goal of the computation?" (e.g. add natural numbers) and "why is it appropriate?" (given the representational convention of the arabic numerals, the algorithm is faithful to the rules governing the operation of addition, which is defined at the level of computation). The level of algorithms and representations, on the other hand, gives a description of the mechanisms that fulfill tasks that are described at the computational level. This level answers "how"-questions ("How are the number representations - numerals – meant to be manipulated"). The level of representations and algorithms describes syntactic or formal means of explicitly representing some of the information needed, and operating upon it "mechanically", whereas the computational level specifies the information represented and operated on.

As Marr stressed, although the algorithms and mechanisms are perhaps empirically easier to access, the computational level is an equally important level from an information-processing theory point of view (Marr, 1982). If we are to explain the workings of a neurocognitive system, the explanation requires as one component a precise understanding of the information extraction problem the system solves. Without the computational level in our theory we can't make head or tails of the neural processes (probably we would not even in principle be able to identify the relevant functional properties of the neural systems except by reference to this higher level).

Of course, there may be more levels of explanation than three that are needed in complete neurocomputational explanations (See for example Craver, 2005; Craver, 2006; Piccinini 2006b). This is not a problem for the point we are making as long as, among those levels, there must be *at least one* that corresponds the Marrian computational level.

With the computational level theory, we can answer many questions about mechanisms. For example, if one considers, *why* this synaptic change is such-and-such, one can answer *because it serves to store the value of x needed in order to compute y*. Or, *why* is the wiring in this ganglion such-and-such? *Because it computes – or approximates computation of – x*. In other words we are able to *explain* many phenomena at the lower levels by their representational character and the appropriateness of the mechanism for the computational task.

In neurocognitive explanation, then, there are *two* angles of attack, two modes of explanation *only one of which is mechanistic*.

One approach is to study the mechanism sustaining the cognitive activity and develop more abstract functional level accounts that enable to state more powerful generalizations about the (counterfactual) behavior of that mechanism. The other is to identify the information extracting problems (such as generation of veridical three-dimensional descriptions from two-dimensional retinal images) as computational tasks, and then *explain the behaviors* in terms of their interpretations. Generally, only with this twofold explanatory capacity is it possible to develop and evaluate an algorithmic theory, since the purpose of algorithmic level is to define algorithms that are used to solve a given task.

This two-pronged approach is, furthermore, characteristic of neuro*cognitive* explanation. Unlike, for example, in the case of the heart, the theory of cognitive competence cannot be simply equated with the totality of (counterfactual) patterns of behavior - the behavioral capacity of the system, if you will. Individual performances always occur in causal interaction with systems that are not part of the neural basis of "the system" as characterized at the level of competence. For example, in the case of a singular visual perception, other neurocognitive mechanisms such as memory mechanisms, attention mechanisms, motivational mechanisms and auditory mechanisms plausibly have an effect on the behavior of the visual system and the outcomes (percepts) produced. These performance factors are abstracted away from in specifying the nature of visual competence.

It makes sense to consider the behavioral capacity as a phenomenon to be explained (cf. Craver, 2006). Indeed, being the set of all possible performances, it is something that needs to be explained. But the explanation can proceed bottom up (from implementing mechanisms) or top-down from the computational theory. This latter is a distinctly different sort of interlevel explanation from the (standard) mechanistic explanation in terms of lower-level implementing causal mechanisms.

In the Marrian model, the higher levels are considered largely autonomous with respect to the levels below, and thus the computational problems of the highest level may be formulated independently of assumptions about the algorithmic or neural mechanisms which perform the computation (Marr, 1982, see also Shapiro, 1997). This amounts to the philosophical position that the level where the cognitive capacities in question are specified is not logically dependent on the ways the causal mechanisms sustain them (Marr, 1982; Shapiro, 1997). As Shapiro (1997) formulates it, the computational level theories are restrictedly regulatory of the algorithmic theory i.e. they are "regulatory", since computational level demands of algorithmic procedures that they produce certain outputs given certain inputs, but since they do not dictate which procedures must be used to produce these outputs, they are "restricted" (Shapiro, 1997).

The main issue is the explanatory role the indispensable and relatively autonomous computational level, and whether the explanations that are based on it are mechanical.

Anderson (1991b, p.471) characterizes the issue thus:

"A rational theory […] provides an explanation at a level of abstraction above specific mechanistic proposals. […] One might take the view […] that we do not need a mechanistic theory, that a rational theory offers a more appropriate explanatory level for behavioral data. This creates an unnecessary dichotomy between alternative levels of explanation, however. It is more reasonable to adopt Marr's view that a rational theory (which he called 'the computational level') helps define the issues in developing a mechanistic theory (which he called the level of 'algorithm and representation'). In particular, a rational theory provides a precise characterization and justification the mechanistic theory should achieve."

We would especially point to two considerations: One, the computational level providing "characterizations and justifications" (answering what and why-questions) is considered a) non-mechanistic, b) above the level of mechanisms, and c) explanatory. Two, computational explanation is seen as consisting of complementary modes of explanation, mechanistic and "rational" (the crucial point here being that in "rational" explanation the computational level is an explanans in it's own right, not a phenomenon to be explained, eventually, mechanistically). These characteristics are in evident and deep-seated conflict with the picture of mechanistic explanation.

Why computational explanation is not (constitutive) mechanistic explanation

Suppose we wanted mechanistic explanations that would include all of Marr's levels, including the computational level (and perhaps many others besides). The level of computation is understood here as a *higher* level of abstraction than that of representation and algorithms. For example, given a computational task, the realization can based on various different representations and various appropriate algorithms might be acceptable for any given representation (which may then be multiply realized in terms of physical implementation).

Therefore explanation where both the computational and algorithmic level are explicitly presented is *multilevel* explanation, since the explanation-relation spans distinct levels of analysis. The computational level entities and principles would have to be either the *explananda* or *explanantia* of this interlevel explanation. Which, then?

Constitutive mechanistic explanation is based on the idea that phenomena at a higher level are always explained by their lower-level constitutive causal mechanisms (Bechtel and Richardson, 1993; Craver, 2001; Craver, 2005; Craver, 2006; Machamer, Darden and Craver, 2000). For example:

"Constitutive explanations are inward and downward looking, looking within the boundaries of X to determine the lower level mechanisms by which the lower level mechanisms by which it can Φ. The explanandum… is the Φ-ing of an X, and the explanans is a description of the organized σ-ing (activities) of Ps (still lower level mechanisms)." (Craver, 2001, p.70).

Being higher in the order of analysis than the level of algorithms, it would seem the competence would then be explanandum, explained by the representation-manipulating mechanisms at the algorithmic (or implementation) level. According to such models competences would always be explained by causal mechanisms operating at the lower level of performances giving answers to how-questions. Indeed, this is what many philosophers seem to have in mind when they talk of mechanistic explanation of the mind. However, in neurocognitive explanations the higher level of organization – viz. computational competence - can be an explanans as well as an explanandum:

"[A] correct explanation of some psychophysical observation must be formulated at the appropriate level. … To be sure, part of the explanation of [the Necker cube's] perceptual reversal must have to do with a bistable neural network (that is, one with two distinct stable states) somewhere inside the brain, but few would feel satisfied by an account that failed to mention the existence of two different but perfectly plausible three dimensional interpretations of this two dimensional image." (Marr, 1982 p.25-26)

It is the interpretations which *explain* what the bistable networks inside the brain are all about. What this also means is that the phenomenon of (mechanisms for) bistable representational states may be explained with reference to the computational task of deriving three-dimensional descriptions of objects from two-dimensional data (and the ensuing ambiguity). Then the object of explanation, the phenomenon, the bistable behavior of the neural mechanism, is at a *lower* level explanation than the interpretation which is doing the explaining.

To take another example, Marr's analysis of the problem of matching parts in two figures which together constitute a random dot stereogram (ibid. pp.111ff) establishes a small set of rules that a stereovision algorithm should respect (and their appropriateness and adequacy, i.e. "rational" grounds for opting to implement them). This analysis then enables the evaluation of the adequacy of presented algorithms (and the inadequacy of some other prima facie plausible algorithms), and therefore a kind of explanation of the behavior of a system running such algorithms (why it uses these particular algoirthms in contrast to the others).

To put things slightly differently, if the explanandum is a performance and the explanation calls for *causal power* for both the mechanism and its constituent elements, then the abstract competence level just cannot serve as a source of explanations. The computational/information structure ("interpretations"), on the other hand, clearly are not characterized in terms of causal organization. It is abstract structure governing (but not causally), or imposed upon, the mechanism. For this reason any talk of "generative mechanisms" operating at the computational level cannot be referring to constitutive mechanisms in the sense the term is used in mechanistic explanation. (So, insofar as they are explanatory, they are explanatory in some other sense of interlevel explanation). It would therefore be a confusion to interpret any talk of "generative mechanisms" as specifications of the computational mechanisms involved in the perfomances, (where "computational mechanisms" are understood as causal, algorithm-level, mechanisms - as discussed for example by Piccinini 2006a; Piccinini 2006b).

Conclusions

A representational system (as opposed to most biological systems) needs to be explained in terms of both the implementing causal mechanisms (or the functional causal organization of the system's performances), as well as an analysis of the representational requirements of the problem solving task or the function represented (the "interpretations" of the functional states qua representational).

The analysis of the competence, i.e. the specification of the system of representation for representing what needs to be represented to get the job done, explains observed performances, i.e. the relevant causal organization.

In computational explanation the abstract competence can serve as explanans, the performances (or behavioral capacities of the mechanism) being the explanandum. But then it is not easy to see how neurocognitive explanation can be accommodated into the constitutive mechanistic model of explanation. In the mechanistic model of explanation the phenomenon to be explained is always at a higher level of explanation, and the explaining is done by giving a description of the lower level mechanism that "sustains" or "produces" the phenomenon. It seems downright impossible to construe such a version of mechanistic model, where the explananda of constitutive explanation *could* lie at lower level.

The problem for extending constitutive mechanistic explanations to computational explanations is that it only recognizes "bottom-up" interlevel explanation – explanation of computation by realizing mechanisms – and thus does can't be reconciled with scientific explanations where the competence level serves as explanans, the performance as the explanandum. If it is the case that in the cognitive sciences theories of competences are used to explain performances, then this mode of explanation is not a case of mechanistic computational explanation of the mind/brain, i.e. a case of explaining the workings of the mind/brain mechanistically. What is explained mechanistically is *how* minds are physically/functionally realized by mechanisms – how it is in principle possible for a physical system to "be a mind". This, however, does not exhaust the notion of computational explanation in neurocognitive sciences.

Based on the fact that in neurocognitive explanations the higher level of organization – viz. computational competence - can be an explanans as well as an explanandum, the character of computational explanation is not captured by the standard forms of mechanistic model of explanation. We concluded that there are two different roles for competences (explanans and explanandum) corresponding to two different modes of explanation in the neurocognitive sciences. Only one of them is covered by the current mechanistic models where explanantia (mechanisms) are located at a lower level of organization than the explananda (phenomena).

Extending the mechanistic model of explanation to cover neurocognitive explanation would mean several major revisions. In current mechanistic models the computational level - a higher level of abstraction - can be seen exclusively as phenomena in need of explanation. This is not in accordance with Marr's original motivation of emphasizing the explanatory relevance of the computational level, where answers to what and why questions are equal in explanatory importance to answers to how-questions.

Thus, there are characteristic and special features of neurocognitive explanation that are not captured by any current version of mechanistic model. This is not to say that the mechanistic model could not conceivably be extended to handle computational competences as explanantia as well as explananda. We are in principle sympathetic to the idea that although the notion of generative mechanisms differs in important explanatory respects from the constitutive causal mechanisms that are explicated in the philosophy of mechanistic explanation, computational-representational explanation could still be interpreted as some kind of mechanistic explanation.

Philosophers have a strong tendency to think of computations as mechanisms (what Marr called algorithms) and the computational hypothesis as a mechanical explanation of the mind. But what is explained mechanically by the computational hypothesis is *how* it is possible for a mechanical system (such as the brain) to be a mind, and exhibit intelligence without causal (mechanical) intervention by an immaterial soul. What is *not* explained mechanically why a brain that was a (specific kind of) mind should be organized thus and so. These explanations are top-down and hence not mechanical.

Acknowledgments

Professor Petri Ylikoski, Professor Matti Sintonen, the members of Philosophy of Science- Research Group (Department of Philosophy, University of Helsinki) and Gualtiero Piccinini.

References

Anderson, J. R. (1991a). The Adaptive Nature of Human Categorization. *Psychological Review*, 98, 409-429.

Anderson, J. R. (1991b). Is Human Cognition Adaptive? *Behavioral and Brain Sciences*, 14, 471-457.

Bechtel, W., & Abrahamsen, A. (2005). Explanation: A Mechanistic Alternative. *Studies in the History and Philosophy of Biomedical Sciences*, 36, 421-441.

Bechtel, W., & Richardson, R. (1993). *Discovering Complexity, Decomposition and Localization as Strategies in Scientific Research.* New Jersey: Princeton University Press.

Chater, N. (1996). Reconciling Simplicity and Likelihood Principles in Perceptual Organization. *Psychological Review*, 103, 566-581.

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic Models of Cognition: Conceptual Foundations. *Trends in Cognitive Sciences*, 10, 287-291.

Churchland, P., & Sejnowski, T. (1992). *The Computational Brain*. Cambridge, MA.: MIT Press.

Craver, C. (2001). Role functions, Mechanisms and Hierarchy. *Philosophy of Science*, 68, 53-74.

Craver, C. (2005). Beyond Reductionism: Mechanisms, Multifield Integration and the Unity of Neuroscience.

Studies in the History and Philosophy of Biomedical Sciences, 36, 373-395.
Craver, C. (2006). When Mechanistic Models Explain. Synthese, forthcoming.
Friston, K. (2002). Functional Integration and Inference in the Brain. *Progress in Neurobiology*, 68, 113-143.
Friston, K. (2005). A Theory of Cortical Responses. *Philosophical Transactions of the Royal Society B*, 360, 815-836.
Glennan, S. (2005). Modeling Mechanisms. *Studies in the History and Philosophy of Biomedical Sciences*, 443-464.
Grush, R. (2001). The Semantic Challenge to Computational Neuroscience. In P. K. Machamer, R. Grush & P. McLaughlin (Eds.), *Theory and Method in the Neurosciences* (pp. 155-172). Pittsburgh: University of Pittsburgh Press.
Machamer, P. K., Darden, L., & Craver, C. (2000). Thinking About Mechanisms. *Philosophy of Science*, 67, 1-25.
Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation of Visual Information*. San Francisco: W.H. Freeman.
Piccinini, G. (2004a). Functionalism, Computationalism and Mental Contents. *Canadian Journal of Philosophy*, 34, 375-410.
Piccinini, G. (2004b). Functionalism, Computationalism, and Mental States. *Studies in the History and Philosophy of Science*, 35, 811-833.
Piccinini, G. (2006a). Computational Explanation and Mechanistic Explanation of Mind. In M. DeCaro, F. Ferretti & M. Marraffa (Eds.), *Cartographies of the Mind: The Interface Between Philosophy and Cognitive Science*. Dordrecht: Kluwer.
Piccinini, G. (2006b). Computational Explanation in Neuroscience. *Synthese*, Forthcoming.
Piccinini, G. (2006c). Computational Modeling vs. Computational Explanation: Is Everything a Turing Machine and Does It Matter for the Philosophy of Mind? *Australasian Journal of Philosophy*, forthcoming.
Shapiro, L. A. (1997). A Clearer Vision. *Philosophy of Science*, 64, 131-153.