KENNETH AIZAWA

REPRESENTATIONS WITHOUT RULES, CONNECTIONISM AND THE SYNTACTIC ARGUMENT[1]

Abstract. Terry Horgan and John Tienson have suggested that connectionism might provide a framework within which to articulate a theory of cognition according to which there are mental representations without rules (RWR) (Horgan and Tienson 1988, 1989, 1991, 1992). In essence, RWR states that cognition involves representations in a language of thought, but that these representations are not manipulated by the sort of rules that have traditionally been posited. In the development of RWR, Horgan and Tienson attempt to forestall a particular line of criticism, the *Syntactic Argument*, which would show RWR to be inconsistent with connectionism. In essence, the argument claims that the node-level rules of connectionist networks, along with the semantic interpretations assigned to patterns of activation, serve to determine a set of representation-level rules incompatible with the RWR conception of cognition. The present paper argues that the Syntactic Argument can be made to show that RWR is inconsistent with connectionism.

In the present paper, I shall argue that Horgan and Tienson have not shown how representations without rules (RWR) could be consistent with connectionism. It seems to me that the Syntactic Argument effectively precludes this. In fact, I believe something stronger is true. It seems to me that RWR is inconsistent with any physical implementation, that nothing in the world could realize RWR. Nevertheless, in order to develop a clear and compelling argument in a limited space, I wish to focus on the consistency of RWR and connectionism. Sections 1.0 and 2.0 will introduce RWR and explain its relationship to connectionism. Section 3.0 will develop the Syntactic Argument in more detail than is given by Horgan and Tienson. After presenting the Syntactic Argument, I examine Horgan and Tienson's two principal strategies for trying to deal with the argument. I argue that Horgan and Tienson's appeal to the fact that representations in connectionist networks can be instantiated in many different patterns of activation cannot defeat the Syntactic Argument (Section 4.1). Nor can their appeal to a concept of *tractability* serve to refute the Syntactic Argument (Section 4.2). In order to simplify the often technically involved presentation of the Syntactic Argument, I ask the reader to hold what will likely be nagging questions concerning tractability until Section 4.2.

Synthese **101**: 465–492, 1994.

1.0. THE THEORY OF REPRESENTATIONS WITHOUT RULES

The representations without rules view may be distilled into three central tenets:

(1) Cognitive systems have syntactically and semantically structured representations (governed by laws/rules/generalizations that are sensitive to the structure and content of the representations).

(2) Mental processing is not governed by so-called quasi-exceptionless, tractable representation-level laws (/rules/generalizations).[2]

(3) Mental processing is governed by soft, tractable representation-level laws (/rules/generalizations).

Understanding RWR requires, most significantly, an explanation of the concept of rule governance, a brief review of the language of thought hypothesis, the three-way distinction between exceptionless, quasi-exceptionless, and soft generalizations, and the concept of tractability. Let me consider these in turn.

1.1. RULE GOVERNANCE AND THE LANGUAGE OF THOUGHT

Claiming that a rule or law or generalization governs the behavior of a system is to implicate it in the causal operation of the system.[3] Although this idea could certainly bear elaboration, for present purposes, it will suffice to note that for Horgan and Tienson the paradigm case of a rule is an instruction in a computer program. For example, in a statement of the more traditional view they oppose, they claim that

> (3) . . . cognitive processing conforms to programmable rules . . . . Claim (3) reflects classicism's assumption that cognitive processing is a matter of there being some (stored or hardwired) *program*. (Horgan and Tienson 1992, p. 28)

The idea of rule governance here must be contrasted with the idea of explicit rule governance. Roughly, computational data structures that represent rules constitute explicit rules, where the instructions of a programming language, such as Turing machine instructions, Pascal instructions, and so forth, constitute what Horgan and Tienson mean by rules *simpliciter*. By this account, Turing machines that compute using the convention that a lone '1' on the tape flanked by 0s denotes one, a lone '11' on the tape denotes two, a lone '111' denotes three, etc. have (are governed by) rules (i.e., Turing machine instructions), but they do not have (are not governed by) explicit rules. The data that appear on the tape represent numbers, not rules. Horgan and Tienson recognize this distinction between explicit and inexplicit rules in the following passage again describing the classical view of cognition they oppose,

> Often, of course, the rules are also contained in the systems as data structures . . . , but that is not necessary . . . . The rules must be general and systematic enough, and interrelated in such a way that they could constitute a program. (Horgan and Tienson 1988, p. 103).

Now, it is not entirely clear to what extent Horgan and Tienson wish to deny the existence or importance of explicit rules in cognition, but it is very clear that they wish to deny the importance of inexplicit rules in cognition.

Aside from a distinction between rules and explicit rules, I should note a distinction between rule governance and rule description. It appears in Horgan and Tienson's writings in a number of forms and a number of places, for example,

> But, we believe, cognitive processes are not driven by or describable by exceptionless rules as required by the standard paradigm. (Horgan and Tienson 1988, p. 97)
> The model is not to contain or *be describable* by such rules. (Horgan and Tienson 1988, p. 103, italics in original)
> Thus, we would like to see connectionist research directed toward structure-sensitive, nonrule-driven, nonsequential processing of structurally rich representations. (Horgan and Tienson 1988, p. 106)
> We also believe that being describable by soft generalizations but not hard rules is characteristic of virtually all of human cognition. (Horgan and Tienson 1988, p. 104, cf. pp. 106, 107,108; Horgan and Tienson 1989, pp. 150, 151,165; Horgan and Tienson 1990, p. 266)

Roughly, the distinction is the following. To describe what a device does is to specify what the device does without commitment to the mechanism that is causally active in the device. A description specifies some behavior without specifying how the behavior comes about. Not every description of what a device does, however, is a description in terms of rules, laws, or generalizations. So, for example, one might describe what a computer, or computer program, does by saying that it keeps track of the company payroll, that it plays chess, or that it simulates a psychotherapist, but these descriptions are not exactly descriptions in terms of rules, laws, or generalizations. To describe a device in these terms is to give a listing of what the device does for each possible input although without commitment to the causal mechanism whereby the device produces its various outputs in response to those inputs. By contrast, if a system is governed by a rule, then that is a fact about the causal mechanisms within.

Now, from the preceding passages, it would seem that Horgan and Tienson would wish to endorse not only theses (1)–(3) above, but also something like,

(1′) Cognitive systems may be said to have syntactically and semantically structured representations.

(2′) Mental processing is not describable by so-called quasi-exceptionless, tractable representation-level laws (/rules/-generalizations).

(3′) Mental processing is describable by soft, tractable representation-level laws (/rules/generalizations).

Against these theses, I will argue that connectionist networks are governed by certain rules, and that they can also be described by these very same rules, so that connectionist networks must be rule describable as well. So, in what follows, I will only directly address claims (1)–(3) in order to overthrow RWR.

Tenet (1) of RWR is intended to be an uncontroversial reaffirmation of the language of thought hypothesis. Some connectionists have challenged this hypothesis, but Horgan and Tienson do not wish to count themselves among their number. Horgan and Tienson believe that mental processing involves simple syntactic items that can be composed to form more complex syntactic items. Further, the complex syntactic items have meanings that are determined by two factors: the meanings of the simple syntactic items that go into them and the way in which the simple syntactic items are combined to form the complex representations. Mental representations are, thus, linguistic in nature; they constitute a sort of language of thought.

1.2. EXCEPTIONLESS, QUASI-EXCEPTIONLESS AND SOFT LAWS

It is common in the philosophy of science, indeed in philosophy in general, to contrast deterministic laws with probabilistic laws. Roughly, a deterministic law is such that, given one state of a system governed by the law, there is only one future development of the system; if the antecedent conditions in a statement of a law obtain, then there is exactly one set of conditions that may obtain as a consequent. A probabilistic law, by contrast, is such that, given one state of a system governed by the law, there are many possible future developments of the system each possibility having a probability given by a probability distribution. In other words, if the antecedent conditions in a statement of a probabilistic law obtain, then one of the consequent states will occur following the probability distribution over the states. This is a familiar taxonomy, but Horgan and Tienson propose to use another less familiar taxonomy in their theory of mind. They propose a three-way distinction between types of laws: *exceptionless* laws, *quasi-exceptionless* laws, and *soft* laws. True laws in the basic sciences, such as physics and chemistry, are frequently taken to be exceptionless. In other words, if the antecedent conditions given in a statement of a true law are satisfied, then so will the consequent conditions. So, for example, supposing that Newton's law of universal gravitation is true, it follows that if any two bodies with masses m_1 and m_2 are separated by a distance r, then they will exert an attractive force on each other equal to Gm_1m_2/r^2, where G is the gravitational constant. As another example, supposing the radioactive decay of uranium atoms is governed by a probabilistic law, there is an exceptionless generalization describing this. If S is a sample of U^{238}, then, with probability P, in time T 50% of S will have undergone radioactive decay. Here we have apparently probabilistic and deterministic exceptionless laws.

In contrast to the laws of the basic sciences, laws in the nonbasic special sciences are commonly assumed to admit of exceptions. Consider a putative psychological generalization such as

(*) If a human wants a beer and believes that she can get one by going to the refrigerator, then she will go to the refrigerator.

This is not a purely exceptionless generalization, for it admits of two types of exceptions. In the first place, this generalization will not describe a person's behavior if she has a massive stroke or is struck by a falling 747 jumbo jet before she reaches the refrigerator. Here the exceptions are not psychological per se. Biological breakdowns, such as a stroke, or external physical interference, such as contact with a crashing 747, are due to contingencies of the biological hardware that happen to constitute a human being or to the contingencies of the physical, nonpsychological interactions a person has with her environment. Some beings that are cognitively identical to humans, but that have an anatomy and physiology differing from humans, will not suffer strokes and will not suffer damage upon being struck by a 747. In addition to the nonpsychological, *inter-level* exceptions to (*), there are psychological-level, representation-level, or *intra-level* exceptions. A human might not go to the refrigerator to get a beer for many *reasons*. She might not want to miss the current conversation, she might not want to drink beer in front of her present company, she might not want to spoil her diet, or she might believe there is a bomb wired to the refrigerator and not want to be blown up. To put the foregoing in other words, we might say that generalizations such as (*) have an implicit *ceteris paribus* clause and that two sorts of exceptions can prevent things covered by (*) from being equal. Things are not equal if there are internal or external physical, chemical, or biological breakdowns and they are not equal if certain psychological background conditions are not in place.

Philosophers of mind generally agree that psychological generalizations admit of inter-level, nonpsychological exceptions due to internal or external physical, chemical and biological factors. *Ceteris paribus* clauses are needed for this reason at least. The issue Horgan and Tienson raise is whether all the intra-level, psychological-level exceptions can be eliminated or not. It is whether the force of the *ceteris paribus* clause must inevitably include intra-level exceptions. According to tradition, scientific investigation will eliminate the psychological-level exceptions to generalizations such as (*), thereby replacing (*) with more precise, accurate and complete generalizations. This will leave psychological generalizations containing only physical, chemical and biological exceptions. It will give psychology quasi-exceptionless generalizations. According to the new RWR conception, however, there are infinitely many psychological-level exceptions to psychological generalizations, hence the exceptions cannot be eliminated. There is simply no way to refine (*) in such a way as to render it absolutely without psychological-level exception. In Horgan and Tienson's terminology, the generalizations in psychology are soft. Thus, where tenet (2) asserts the negative thesis that psychological processes are not governed by quasi-exceptionless representation-level generalizations, tenet (3) asserts the positive thesis that psychological processes are governed by soft representation-level laws. Note that although RWR stands for representations without rules, this does not mean that there are not supposed to be any rules in cognition at all. It only means that there are no quasi-exceptionless rules of the sort cognitive science commonly assumes.

For the sake of clarity, it is worth stating the distinction between exceptionless, quasi-exceptionless and soft laws in such as way as to facilitate contrasting them with deterministic and probabilistic laws. A law is exceptionless if, given the antecedent conditions obtain, the consequent condition obtains. A law is quasi-exceptionless if, given the antecedent conditions and some fixed background of 'lower level' conditions, the consequent condition obtains. A law is soft if, even given the antecedent conditions and some fixed background of 'lower level' conditions, no specified consequent condition of the law need obtain.

The rationale for rejecting the traditional view that quasi-exceptionless rules will ultimately be found to govern psychological processes is simple. It is merely a different extrapolation from examples than is usually made. Horgan and Tienson discuss at length the factors that can go into the determination of a professional basketball player's choice of action, say, to bounce pass to a teammate or to take a jump shot. These factors include the speed at which a player is running, how well the player is shooting, how well the player is passing, the speeds and positions of his four teammates, how well they are shooting, how well they are catching passes, how well they are passing, the speeds and positions of the five opponents, their defensive strengths and weaknesses, the positions of the referees, the time remaining in the game, the tide and tempo of the game and on and on (cf. Horgan and Tienson 1989, pp. 97–102). After developing this long list, they suggest that such lists are in fact typical of what is involved in most human decision making. They propose that the list of relevant variables cannot, even in principle, be completely specified.

1.3. TRACTABILITY

A final crucial element of RWR is the concept of tractability. Although Horgan and Tienson mention tractability, or something akin to it, in most of their presentations of RWR, they have not developed the concept in any detail. For example, in 1988, they wrote, Fifth, the rules are *tractable*. The rules must be general and systematic enough, and interrelated in such a way that they could constitute a [computer] program. Though it is hard to say exactly what this amounts to, the intuitive idea seems clear enough. A mere list of all possible inputs with the resulting output for each, for instance, would not qualify as a set of tractable rules. (Horgan and Tienson 1988, p. 103; Horgan and Tienson 1989, p. 163, n. 19; Horgan and Tienson 1990, p. 260)

The concept of tractability appears as something of an afterthought in Horgan and Tienson's development of RWR. In their first paper, it is introduced in one paragraph only to be set aside in the next (Horgan and Tienson 1988, p. 103). In their second paper, it is mentioned only in a footnote. In their third paper it is no longer mentioned by name. In their latest paper (Horgan and Tienson 1992), the concept does not appear at all. I bring the concept to the fore in my presentation of RWR primarily because it will play a prominent part in Horgan and Tienson's attempt to respond to the Syntactic Argument. Note, however, that the idea of tractability is logically independent of the idea of a law's being exceptionless, quasi-exceptionless, or soft. There can, for example, be intractable exceptionless, quasi-exceptionless and soft laws.

2.0. RWR AND CONNECTIONISM

Although Horgan and Tienson mean for the RWR conception of cognition to stand or fall on its own merits, they have nonetheless, developed it more or less through the guidance of the Parallel Distributed Processing brand of connectionism (McClelland, Rumelhart and the PDP Research Group 1986). In their 1988 paper, Horgan and Tienson claimed that existing connectionist networks failed to instantiate RWR because those networks did not use structured representations, hence could not satisfy the first tenet. In this paper, they wrote,

To our knowledge, none of the work done to date in the connectionist framework falls within this [RWR] region. But we believe there is reason to hope that a version of connectionism can occupy this region. (Horgan and Tienson 1988, p. 97)

Then, when explaining their reasons for claiming that connectionist work to date had not illustrated the RWR conception, they reviewed reasons for the view that existing connectionist networks do not use structured representations, hence do not satisfy tenet (1). Horgan and Tienson nevertheless expressed the hope that Paul Smolensky's tensor product theory of representation (Smolensky 1989) might show how compositional representations are possible in connectionist networks, hence how connectionist networks might satisfy all the tenets of RWR (Horgan and Tienson 1988, p. 106). By 1989, Horgan and Tienson were apparently more confident of Smolensky's theory, claiming that tensor product theory provides a "rich and robust notion of constituent structure, and thereby the basis for compositional syntax" (Horgan and Tienson 1989, p. 165).[4] Hence, they concluded that existing connectionist models can instantiate RWR: "There are links between RWR and connectionism: connectionism shows one way that RWR might be realized" (Horgan and Tienson 1989, p. 168, cf. the more cautious statement on p. 148). In personal correspondence, Horgan has told me that the "might" in the preceding passage should be interpreted in an epistemic fashion and that they take it as an open question whether connectionist RWR is possible. Horgan and Tienson make this latter point in their most recent installment on RWR, "It is an open question, at present, whether or not there can be *connectionist* cognitive systems with the features that characterize the RWR conception of cognition" (Horgan and Tienson 1992, p. 29). With this in mind, one may construe the present paper as an attempt to press the Syntactic Argument for the conclusion that connectionist RWR processing is not possible.

In asserting no more than that connectionism shows one way RWR (epistemically) might be realized, it is not clear whether Horgan and Tienson mean to make the strong claim that perhaps *every* connectionist model instantiates the RWR conception or merely that perhaps there are some connectionist models that instantiate RWR. In general, I believe that Horgan and Tienson mean only the weaker claim. Be this as it may, however, I propose not to do an exegetical analysis of the matter, since I believe that even the weaker version can be challenged using the Syntactic Argument.

3.0. THE CONSISTENCY OF CONNECTIONISM AND RWR: THE SYNTACTIC ARGUMENT

The philosophical literature on connectionism contains reason to believe that there is some tension between connectionism and RWR. For example, in 1988, Fodor and Pylyshyn argued that connectionist networks qua connectionist networks cannot instantiate the sort of compositional representations endorsed in tenet (1) of RWR (Fodor and Pylyshyn 1988). Only connectionist networks qua implementations of Classical computer architectures can do this. Fodor and Brian McLaughlin pursued this line of argumentation showing that Smolensky's tensor product theory will not produce the syntactic and semantic compositionality they claim is required to account for the productivity and systematicity of thought (Fodor and McLaughlin 1990). The Syntactic Argument supports a stronger conclusion than that emerging from the Fodor–McLaughlin–Pylyshyn line. Rather than merely arguing that connectionist models qua connectionist models cannot instantiate RWR, the Syntactic Argument shows that no available connectionist models at all can instantiate RWR. It shows that *all* existing networks are governed by quasi-exceptionless representation-level rules. Horgan and Tienson have foreseen the Syntactic Argument and have attempted to forestall it in their development of RWR. Their first response involves an appeal to the fact that connectionist representations might be instantiated in many distinct patterns of activation. Waiting in the wings is a response based on a concept of tractability. Unfortunately for RWR, neither of these defenses is cogent.

While the Fodor–McLaughlin–Pylyshyn argument is directed against tenet (1), the Syntactic Argument attacks tenets (2) and (3). Here is how Horgan and Tienson present it,

> It is indisputable that a connectionist network is rule-characterizable at *some* level of description, viz., at the level of the individual nodes, whose computations and interactions conform to precise mathematical algorithms. Call these *node-level* rules (NL rules). And if there are structured representations in a connectionist model, there must be precise rules specifying how both the atomic constituents of these representations and their combinatorial syntax are realized in the system. Call these *representation instantiation* rules (RI rules).
>
> But if there are hard NL rules describing the behavior of the nodes, and hard RI rules determining the representational structure of the systems from its node level description, then it would seem that it must *follow* that the system will be describable by hard rules that advert entirely to the compositional structure of the representations. That is, it appears, representation level rules will simply follow from the NL and RI rules. Hence, an RWR version of connectionism is impossible. (Horgan and Tienson 1989, p. 106)

Notice that the conclusion of the argument is that the system with NL rules and RI rules will be describable by hard [i.e., quasi-exceptionless] rules. As mentioned above, I will argue against this conclusion by showing that a network is governed by a set of quasi-exceptionless rules. These rules can then serve as a description of the network.

Given the centrality of the Syntactic Argument to the present attack on RWR, it is practically essential that I elaborate upon it through the use of examples. To make the situation as clear as possible, I begin with an extremely simple network showing how the Syntactic Argument works. From this it should be much clearer how more complex networks are handled. I think that most of the philosophical work will be done in the first simple case with the remaining cases clearly following from mere mention of the technical details. Because of this, the development of this first simple case will go rather slowly beginning with some slightly technical details, then turning to the philosophical interpretation. For simplification, I shall ignore the issue of tractability in this section and return to it in Section 4.2.

Consider a deterministic four-node network with input states determined by two nodes and output states determined by two nodes. A network of this type could be a simple feedforward network with two input nodes and two output nodes or an asynchronous update network, such as a Boltzmann machine or a Hopfield net. Suppose that the input and output states are given by strings of 1s and 0s, so that however the nodes are interconnected and whatever the weights on the connections, the network in question will compute one of the 256 possible mappings from pairs of 0s and 1s to pairs of 0s and 1s. At times it will be convenient to use an asterisk as a variable for the value of a node's activation. Thus, $00 \rightarrow **$ means that the input activation pattern 00 is mapped onto a pair of activation values, although it does not matter which.

Horgan and Tienson allow that input and output patterns of activation have semantic interpretations given by the so-called representation-instantiation rules (RI rules). This is how connectionist networks are supposed to satisfy tenet (1) of RWR. One way of generating representation-level rules for a simple four-node network with two input nodes and two output nodes is to pair each of the possible inputs described representationally with its output described representationally. Schematically, this yields quasi-exceptionless representation-level rules of the form, (I),

If $RI_i(00)$, then $RI_o(**)$
If $RI_i(01)$, then $RI_o(**)$
If $RI_i(10)$, then $RI_o(**)$
If $RI_i(11)$, then $RI_o(**)$,

where the subscripts indicate whether a representation instantiation rule applies to an input pattern of activation or an output pattern of activation. These rules might be read, if the cognizer is in the input psychological state $RI_i(**)$, then it will go into the output psychological state $RI_o(**)$.[5] As (I) now stands it is not entirely adequate. Three conventions that will not be reflected in the orthography of (I) must be installed. First, and most simply, if some input pattern of activation has, for some reason, no semantic interpretation, then there is no corresponding rule. Second, if two input patterns of activation receive the same semantic interpretation, e.g., if $RI_i(00) = RI_i(01)$, then the rules instantiating the form must be modified. If two rules are identical, having both the same antecedent and the same consequent, then one rule is simply eliminated. If, however, two rules have the same antecedent, but distinct consequents, then the two rules are combined into one rule. The new rule will have the same antecedent as the old rules and the consequent of the new rule will be a disjunction of the consequents of the old rules. Obviously, the same convention holds when more than two rules have the same antecedents. Third, during any network computation, the rules instantiating the rule form (I) apply in parallel, rather than in serial. Whichever rule, if any, has its antecedent condition satisfied is applied. This captures the parallel processing in connectionist networks. Of course, these rules might be implemented serially, but that is mere implementation. In reality, the rules in a four-node, two-layer network apply in parallel.

To illustrate what this notation for rule forms means, consider three examples of ways in which these rule forms might be instantiated. Imagine a network that computes the following node-level input-output mapping:

$$00 \rightarrow 00$$
$$01 \rightarrow 11$$
$$10 \rightarrow 11$$
$$11 \rightarrow 11,$$

and let the RI_i and RI_o rules be the following,

$RI_i(00)$ = no dog is present
$RI_i(01)$ = a dog is present on the right
$RI_i(10)$ = a dog is present on the left
$RI_i(11)$ = dogs are present on the right and the left

$RI_o(00)$ = do not start fleeing
$RI_o(01)$ = start fleeing

$RI_o(10)$ = start fleeing
$RI_o(11)$ = start fleeing.

Then, the rule forms are instantiated as follow

If no dog is present, do not start fleeing,
If a dog is present on the right, start fleeing,
If a dog is present on the left, start fleeing,
If dogs are present on the right and the left, start fleeing.

More explicitly, the first rule would be read, if the cognizer is in the input psychological state "no dog is present", then it will go into the output psychological state "do not start fleeing". In this example, no reduction in the number of rules is necessary and none is implied by the notation. As a second example, keep the same input-output mapping at the node level and keep the same RI_o function, but change the RI_i function to the following,

$RI_i(00)$ = no dog is present
$RI_i(01)$ = a dog is present
$RI_i(10)$ = a dog is present
$RI_i(11)$ = a dog is present.

In this case, the forms are instantiated by

If no dog is present, do not start fleeing,
If a dog is present, start fleeing,
If a dog is present, start fleeing,
If a dog is present, start fleeing.

In this example, elimination of redundant rules would leave us with only two rules for our network. As a third example, again leave the node-level input-output mapping and RI_o the same, but change $RI_i(01)$ to yield,

$RI_i(00)$ = no dog is present
$RI_i(01)$ = no dog is present
$RI_i(10)$ = a dog is present
$RI_i(11)$ = a dog is present.

In this case, the representation-level rules become

If no dog is present, do not start fleeing,
If no dog is present, start fleeing,

If a dog is present, start fleeing,
If a dog is present, start fleeing.

Here the first two rules may be combined and the fourth rule removed, since it is, redundant. This leaves,

If no dog is present, start fleeing or do not start fleeing,
If a dog is present, start fleeing.

In all explicitness, the first rule should be read "If in an input psychological state 'no dog is present', then go into the output psychological state 'start fleeing' or into the output psychological state 'do not start fleeing'". The combination of the first two rules yields a disjunction, rather than a conjunction, since it is never the case that the network simultaneously does, or thinks to do, both. Instead, the network will do (or think to do) one or the other, depending on the pattern of activation making up the input. Note as well that this last set of rules is not deterministic: a given input state does not lead to a unique output state. This, however, does not matter to the issue of quasi-exceptionless. It may still be the case that if the antecedent of the law is satisfied, then if all things at the psychological level are equal, then the consequent will be satisfied as well. This is what it is for a law to be quasi-exceptionless.

Using the material developed to this point, I wish to argue that no four-node network with two input nodes and two output nodes having binary-state nodes instantiates RWR. It should be clear that for any network of this type, there will be some set of rules of the form (I), i.e., of the form

If $RI_i(00)$, then $RI_o(**)$
If $RI_i(01)$, then $RI_o(**)$
If $RI_i(10)$, then $RI_o(**)$
If $RI_i(11)$, then $RI_o(**)$.

These rules constitute quasi-exceptionless rules governing the behavior of the sort of simple network now under consideration. Barring exceptions such as a stroke or being run over by a car, a simple organism with nothing more for a brain than the sort of four-node network just introduced will be governed by some set of laws of this form. We will have to address the question of tractability at some point, but we will come to that.

Here I can certainly imagine some resistance to asserting that rules such as "If in an input psychological state 'no dog is present', then go into the output psychological state 'start fleeing' or into the output psychological state 'do not start fleeing'" govern the behavior of the simple networks I have described. Although I can see no *reason* for this resistance and cannot possibly foresee and forestall every possible line of resistance, I do wish to make a simple observation that will make it hard to say that my rules do not govern networks, while at the same time maintaining that soft rules govern networks. To defend RWR by attacking my assumption about the causal role of these rules, one cannot simply provide reasons for thinking that networks are not governed by the rules I say they are. One must put forth a sense of rule governance that can reject my proposal concerning which rules are at work, while not rejecting the RWR proposal concerning which rules are at work. To save RWR with a theory of what rules truly govern the networks, one must show how soft rules, but not quasi-exceptionless rules, may be said to govern networks. That, I think, is an especially difficult challenge.

I take it that I have now secured the claim that RWR cannot be instantiated in the sort of simple four-node net described above. This, of course, does not show what I ultimately wish to claim, namely, that no connectionist network at all can instantiate RWR. Let me consider, then, various complications. Consider a four-node network of probabilistic, two-state nodes. The NL rules in such a net will be of the form,

If 00, then
  00 with probability p^{00}_{00},
  01 with probability p^{00}_{01},
  10 with probability p^{00}_{10},
  11 with probability p^{00}_{11},
If 01, then
  00 with probability p^{01}_{00},
  01 with probability p^{01}_{01},
  10 with probability p^{01}_{10},
  11 with probability p^{01}_{11},
If 10, then
  00 with probability p^{10}_{00},
  01 with probability p^{10}_{01},
  10 with probability p^{10}_{10},

11 with probability p_{11}^{10},
If 11, then
00 with probability p_{00}^{11},
01 with probability p_{01}^{11},
10 with probability p_{10}^{11},
11 with probability p_{11}^{11},

where the superscripts on the probabilities indicate the relevant inputs and the subscripts indicate the relevant outputs, and where

$$p_{00}^{00} + p_{01}^{00} + p_{10}^{00} + p_{11}^{00} = 1,$$
$$p_{00}^{01} + p_{01}^{01} + p_{10}^{01} + p_{11}^{01} = 1,$$
$$p_{00}^{10} + p_{01}^{10} + p_{10}^{10} + p_{11}^{10} = 1,$$
$$p_{00}^{11} + p_{01}^{11} + p_{10}^{11} + p_{11}^{11} = 1.$$

At this point the representation instantiation rules may be applied with the values of RI_i(**) in the antecedents and the values of RI_o(**) in the four slots in the consequents. This yields representation-level rules of the form, (II),

If $RI_i(00)$, then
$RI_o(00)$ with probability p_{00}^{00},
$RI_o(01)$ with probability p_{01}^{00},
$RI_o(10)$ with probability p_{10}^{00},
$RI_o(11)$ with probability p_{11}^{00}.
If $RI_i(01)$, then
$RI_o(00)$ with probability p_{00}^{01},
$RI_o(01)$ with probability p_{01}^{01},
$RI_o(10)$ with probability p_{10}^{01},
$RI_o(11)$ with probability p_{11}^{01}.
If $RI_i(10)$, then
$RI_o(00)$ with probability p_{00}^{10},
$RI_o(01)$ with probability p_{01}^{10},
$RI_o(10)$ with probability p_{10}^{10},
$RI_o(11)$ with probability p_{11}^{10}.
If $RI_i(11)$, then
$RI_o(00)$ with probability p_{00}^{11},
$RI_o(01)$ with probability p_{01}^{11},
$RI_o(10)$ with probability p_{10}^{11},
$RI_o(11)$ with probability p_{11}^{11}.

Again, we want the rules of form (II) to respect the earlier conventions governing rules of form (I). Input patterns of activation that have no semantic interpretation correspond to no rules. If two input patterns of activation are assigned the same semantic interpretation, we combine rules in the former way, modifying the probability distribution over the output states in the obvious way. Contrary to RWR, these rules also govern the behavior of such networks. This suffices to show that no two-state, probabilistic networks with four nodes instantiates RWR. The only philosophical point that might bear further emphasis here is that even though the rules instantiating form (II) are only probabilistic, they are also quasi-exceptionless. If the antecedent conditions of the rules are satisfied and there is no implementation level interference, then the conditions of the consequent will be realized. One of the events in the consequent will occur in accordance with the appropriate probabilities.

Here I take it as obvious how to extend the above schemata to include what we might call 'random' two-state networks where there is no probability distribution on the consequent states. Simply exclude the references to probabilities. I think it is also obvious how to incorporate k-state, rather than mere two-state nodes.[6] This secures the claim that no four-node network at all can instantiate RWR, unless the tractability condition can be made to rescue RWR. To establish the claim more generally, I must consider networks with more nodes. It should be clear how to extend the rule schemata for networks with more nodes in the input or output layers of the network so rehearsing the argument for them is really unnecessary. This leaves the case in which hidden nodes are added to a network. Suppose, then, that one adds, say, two hidden nodes to the four-node network. These hidden nodes may have a semantic interpretation or not. If they do not, then they may be relegated to the status of implementation detail and ignored; the representation-level rules are unchanged by the introduction of any number of uninterpreted hidden nodes. If the hidden nodes do have a semantic interpretation, then we must recognize two sets of representation-level rules. The first set governs the transition from the representations on the input nodes to the representations on the hidden nodes, while the second set governs the transitions from the representations on the hidden nodes to the representations on the output nodes. The point, stated more generally, is that we need to postulate a set of representation-level rules for all and only the network mechanisms that map representations to representations. Several sets of representation-level rules applied serially govern the behavior of some networks.

This last general point turns out to be crucial to a proper understanding of the way the Syntactic Argument must be run against RWR. Suppose the Syntactic Argument were not understood this way. It might then be objected that the Syntactic Argument is too strong to be correct. If the Syntactic Argument formulated above were a sound argument for showing what rules govern a process, then all cognitive processes would be governed by mere look-up tables, but that cannot be correct. Consider the play of chess. One may suppose, crudely, that the input to some chess-playing center in the brain is a mental representation of the layout of the pieces on the 8×8 chessboard and that the output is another mental representation of the layout of the pieces on the board. If the Syntactic Argument were correct, then we could say that the mind of the chess player is governed by a set of rules that maps the very large number of possible chess positions to the very large number of possible responses to those positions. This conception of the cognitive activity of chess playing is certainly mistaken; chess playing is evidently not governed by a 'look-up table'. So, any version of the Syntactic Argument that leads to the conclusion that the chess player is governed by such laws must be fallacious.

The above is a misapplication of the Syntactic Argument, at least in the way I wish to develop it. The Syntactic Argument does not lead to the view that the mind is one large look-up table. Instead, it maintains that the connectionist networks leading from representational state to representational state constitute look-up tables. The rules generated by the Syntactic Argument are rules that connect one representation to the next, when no further representations intervene. To omit intervening representations would be to omit intervening cognitive steps. To put this in another manner, the Syntactic Argument does not involve differing with Horgan and Tienson on exactly what representational states appear in cognitive processing. Let Horgan and Tienson specify the inputs to and outputs from cognitive processing, as well as all the intermediate representations (this they must do in accordance with tenet (1)) and let them specify the type of network they take to be at work 'beneath' these representations. The Syntactic Argument then shows that, given these representations and the node-level rules underlying them, there must be some quasi-exceptionless representation-level rules mediating the transitions from one set of representations to another.

4.1. THE MULTIPLE REALIZABILITY OF REPRESENTATIONS DEFENSE

Horgan and Tienson anticipated the Syntactic Argument in their first installment on the RWR view in 1988 and therefore had a lot to say following their presentation of it. Most of their comments are true, but none of them genuinely addresses the Syntactic Argument. Here is the main thrust of their response:

> Suppose you have a simple connectionist network trained up so that if it is given input A alone, it goes into output state C. And if it is given input B alone, it goes into state D. But states C and D are incompatible. This system is then correctly describable by the *ceteris paribus* (CP) generalizations, "If A, then CP, C", and "If B, then CP, D".
>
> What if we now give the system input A and B together? We do not know, of course. We do not have enough information. But the interesting fact is, given only that the input is A + B, the system could be such that there is no determinate fact of the matter about what output state it will settle into . . . .
>
> Here are some of the simple reasons why this can happen.
>
> (1) Suppose, for example, identical networks have been trained up from different small random weights (as is normal), so that both obey the simple generalizations, if A, then CP, C and if B, then CP, D. Having started from different weights, they will have different final weights. Because of this, it might happen that when given input A + B one settles in C and the other settles in D. (Horgan and Tienson 1988, p. 107)

Here I might interject an expository point. John Tienson has told me in personal communication that we are to understand their view in the strong fashion: they wish to assert that there may be no fact of the matter as to what the net will do on A + B.

This response will not suffice to defend RWR. There is some fact of the matter as to what the net will do. Recall that the Syntactic Argument began by showing that (issues of tractability aside) *for any* four-node deterministic network of two-state nodes, there exists some set of representation-level rules of the form,

If $RI_i(00)$, then $RI_o(**)$
If $RI_i(01)$, then $RI_o(**)$
If $RI_i(10)$, then $RI_o(**)$
If $RI_i(11)$, then $RI_o(**)$,

governing the operations of these networks. From this case, we generalized the argument to larger networks to show that, setting the issue of tractability aside, no connectionist network can instantiate RWR. Now, in response, Horgan and Tienson note that a given representational state may be implemented by two distinct [input] patterns of activation and that these distinct patterns of activation may result in different representation-level outputs. But, recall that, throughout the development of the Syntactic Argument, I was at pains to point out how the rule forms have certain implicit conventions for simplification when a single input representation might lead to distinct output representations. This complexity of exposition was required in order to handle easily this response by Horgan and Tienson. Nothing is affected by the multiple instantiability of representations in distinct patterns of activation.

4.2. THE TRACTABILITY DEFENSE

So far I have run the Syntactic Argument putting off the question about the tractability of the rules that are to be generated by the Syntactic Argument. Now I must face this issue head on. I have two responses to the tractability condition. In the first place, even if it is the case that look-up table rules of the sort I have described do not count as tractable, this does not save what is supposed to be the most important idea of RWR, namely, the idea that cognition involves soft laws. The most that this tractability response can do in the face of the Syntactic Argument is show that connectionist networks are governed by intractable, quasi-exceptionless laws. This is far less than Horgan and Tienson had hoped to show. It hardly heralds the introduction of a new paradigm in cognitive science. Be this as it may, it seems to me that Horgan and Tienson cannot make stick the claim that the exhaustive lists I have given are intractable. I contend that Horgan and Tienson give no conception of tractability that is strong enough to rule out the sorts of rules that I have argued govern connectionist networks. I shall proceed by first reviewing what Horgan and Tienson have to say regarding tractability, explaining why what they propose does not exclude the sets of rules proposed in my version of the Syntactic Argument. I will then provide a similar treatment of other proposals.

In the passage from Horgan and Tienson's (1988) paper where they first mentioned tractability and its rationale, they suggested that a necessary condition on a set of rules being tractable is that it have the form of a computer program. This condition, however, is not strong enough to exclude the rules generated by my version of the Syntactic Argument. The rules the Syntactic Argument postulates meet the necessary condition on tractability. To see this, we might note that there are simple Pascal programs that are essentially mere lists specifying outputs for all possible inputs. For example, let a variable 'test_result' have as possible values elements in the set {0, 1, 2, 3}. The program is

```
case test_result of
   0: write(0);
   1: write(1);
   2: write(2);
   3: write(3);
   end;
```

Here a single instruction codes a look-up table of what is to be done for each possible value of the variable 'test_result'. It appears that any CASE-statement with a finite number of cases will count as a finite look-up table and a legitimate computer program. Further, a sequence of two or more CASE-statements will also be a legitimate Pascal program. The rules that I suggest govern a connectionist network implementing a transition between representations are evidently just like the CASE-statements of Pascal.

I might reinforce this intuitive example with a brief glance at the foundations of computation theory. It shows that tractability defined as computability does not rule out the sorts of rules I have described. In a classic text on recursion theory, Rogers gives the following five conditions on the intuitive notion of an algorithm:

*1. An algorithm is given as a set of instructions of finite size. (Any classical mathematical algorithm, for example, can be described in a finite number of English words.)

*2. There is a computing agent, usually human, which can react to the instructions and carry out the computations.

*3. There are facilities for making, storing and retrieving steps in a computation.

*4. Let P be a set of instructions as in *1 and L be a computing agent as in *2. Then L reacts to P in such a way that, for any given input, the computation is carried out in a discrete stepwise fashion, without use of continuous methods or analogue devices.

*5. L reacts to P in such a way that a computation is carried forward deterministically, without resort to random methods or devices, e.g., dice (Rogers 1987, p. 2).[7]

As additional conditions, he proposes the following (which have been paraphrased):

*6. There is no fixed finite bound on the size of inputs.

*7. There is no fixed finite bound on the size of a set of instructions.

*8. There is no fixed finite bound on the amount of 'memory' storage space available.

*9. There is a fixed finite bound on the capacity or ability of the computing agent

*10. There is no bound on the length of a computation (cf. Rogers 1987, pp. 3–5).

*1 asserts the finiteness of programs, yet nothing in this list is anything like Horgan and Tienson's concept of tractability. Other treatments of computability that are less explicitly devoted to foundational detail apparently embody Rogers' conception as well. Davis (1982) for example, gives the following definition of a Turing machine: "A **Turing machine** is a finite (nonempty) set of quadruples that contains no two quadruples whose first two symbols are the same" (Davis 1982, p. 5). In another text, Cutland does not discuss anything like Horgan and Tienson's conception of tractability, nor does he give an extensive discussion of the intuitive concept of effective computability or of an algorithm, but he does require that programs for so-called *unlimited register machines* contain only finitely many instructions, that programs for Turing machines contain only finitely many instructions, and that *Post production systems* have only finitely many productions (Cutland 1980, pp. 9, 54, 59).[8] Machtey and Young (1978) proceed as does Cutland. They do not mention anything like Horgan and Tienson's tractability requirement, but develop computational formalisms having the finiteness condition. They require that so-called *random access machine* (RAM) programs, Turing machine programs and *Markov algorithms* contain only a finite number of instructions (Machtey and Young 1978, pp. 28f, 33f, 38f).[9,10]

I think the foregoing considerations show that the rules generated in my version of the Syntactic Argument are tractable in the sense of being computable by a Turing-equivalent device. Still, it is worth surveying other possible senses of tractability that might eliminate the sets of rules I have suggested govern networks. Perhaps Horgan and Tienson might relate tractability to concepts in the theory of computational complexity. For present purposes, I must ignore a great number of small qualifications that would make the exposition technically correct on all points. What I ignore will, of course, not affect the ultimate outcome of the argument. For a more elaborate development of complexity theory than is possible here, see Machtey and Young (1978, Chap. 5) and especially, Garey and Johnson (1979). Measures of computational complexity are meant to capture the intuitive idea of how difficult it is to perform some computation. Computational complexity measures are (1) functions of input size, (2) they specify some sense of 'how hard' it is to compute some function and (3) they make this specification relative to some specific computing device. Accepting (3), there must be complexity measures for Turing machine programs, for random access machine programs, for unlimited register machine programs, and so on. Consider Turing machines. To measure input sizes for Turing machines, following (1), one might most naturally count the number of tape squares the input occupies. Since not every computing device has a tape of the sort found in Turing machines, this measure of input size is defined relative to Turing machines. One reason for making complexity measures functions of input size is that this enables us to capture the intuitive idea that it is more difficult to compute the values of a function on larger inputs. That is, it is intuitively harder to square 12,345 than it is to square 2, it is harder to add 543,245 and 3,445 than it is to add 1 and 1. Having fixed upon a specific computing device and a particular measure of input size, there are still various standards by which to assess the difficulty of computing a particular function. For example, there is the maximal number of tape squares used in the course of any computation of the function f on any input of size n or the maximal number of instructions executed in the course of any computation of the function f on any input of size n. The first of these measures might be called the *space complexity* of a program for a function, where the second is the *time complexity* of a program for a function. By specifying (2)–(3), we obtain a number of precise ways of stating how hard it is for some machine to compute some function. Most purposes in computation theory and AI are served by recognizing two principal categories in which to place the functions defining complexity measures. Complexity measures that increase as a polynomial function of the size of inputs, for example, $C(n) = n^2$ or $C(n) = n^3$, are typically counted as tractable, where complexity measures that increase as an exponential function of input size, for example, $C(n) = 2^n$ or $C(n) = 3^n$ are taken to be intractable.

We are now in a position to see, in at least a sketchy fashion, the extent to which computational complexity theory supports Horgan and Tienson's claim that the set of rules the Syntactic Argument sets forth are, in fact, intractable. As things stand now, it is unclear how the resources of complexity theory can be brought to bear on connectionism. Too many conceptual questions remain to be answered. To begin with, some measure of input size for connectionist networks must be defined. This seems simple enough; simply count the number of nodes used in the representation of the inputs. Next, the claim that exhaustive lists are intractable must be interpreted as a claim about the rate of increase of the number of rules or instructions in the program as the size of the inputs increases. Here is the conceptual problem. A typical connectionist network only handles inputs of a fixed size n determined by the fixed number of input nodes in the net. It, therefore, makes no sense to say either that the number of rules needed increases either polynomially or exponentially. To put it more technically, the function that takes input sizes as arguments and gives as values the number of rules it takes a given network to compute the squaring function is defined at only one point. That point is equal to the number of input nodes in the net. One point is not sufficient for determining a function to be either exponential or polynomial, hence either tractable or intractable. The obvious problem, then, is that there is no basis for saying that the number of instructions a given network uses is either a polynomial or exponential function of the input size, hence no means of applying standard complexity theory.

The preceding passes through computation theory and complexity theory may well be overly technical. Perhaps what is intended is something less technical and more intuitive. Earlier we said that the closest thing standard computation theory appears to offer in the way of a prohibition on the sorts of look-up tables I have postulated is the requirement that they be finite. There is no fixed finite size on them, but they must be finite. Standard computation theory prohibits infinite, but not finite look-up tables. Perhaps tractability might be interpreted a placing a restriction on the size of finite look-up tables. Is there not some sense in which a look-up table is intractable if it has more instructions than there are neurons in the brain? Is there not some sense in which a look-up table is intractable if its implementation in the brain takes longer to run that one human lifetime?[11] The answers to both of these questions is 'yes', but this does little to blunt the force of the Syntactic Argument.

Recall the nature of the Syntactic Argument. It showed that for any network Horgan and Tienson might propose as a theory of human cognition, there exists a set of quasi-exceptionless, representation-level rules that govern the network's behavior. Thus, let Horgan and Tienson specify a network that satisfies their demands for limitations on the time of computation and the size of the network. The Syntactic Argument as I have developed it simply shows how to find the quasi-exceptionless representation-level rules governing the network that Horgan and Tienson provide. The argument does not involve adding more layers to the network, thereby increasing the amount of time it takes the network to run. The argument does not add any nodes at all, so that there are no increased demands on the amount of material in the brain. So, whatever limitations on time and material resources Horgan and Tienson wish to have respected in their network theory of the brain can be respected in the network theory constructed using the Syntactic Argument.

The upshot of this section is that it does not really matter for Horgan and Tienson's theory of soft laws whether or not the look-up table rules are tractable, since in any case the hypothesis that soft-law-driven cognition can be implemented in connectionist networks is false. This is enough to undermine the central idea of RWR. Moreover, it turns out that Horgan and Tienson do not have available a concept of tractability that would protect their preferred version of RWR (one that does not use analogue representations) against the Syntactic Argument.

5.0. CONCLUSION

The central claim of the present paper is that the Syntactic Argument in fact proves that the central idea of RWR cannot be instantiated in connectionist networks. That is, it appears that the combination of node-level laws with representation-instantiation laws determines a set of quasi-exceptionless, representation-level laws governing the behavior of connectionist networks. Further, Horgan and Tienson do not have a concept of tractability that will defend their version of RWR against the Syntactic Argument. Although there may yet be means by which the RWR conception can be rendered consistent with connectionism, a re-evaluation of Horgan and Tienson's responses to the Syntactic Argument would seem to be an appropriate place with which to begin examining the viability of RWR.

NOTES

[1] The present paper has been improved by conversations with Terry Horgan and John Tienson. Thanks are also due to Gary Fuller, John Heil, Terry Horgan and Bob Stecker for comments on earlier drafts of this paper.

[2] Horgan and Tienson use the terms 'law', 'rule', and 'generalization' interchangeably. This, despite the fact that, for example, laws are counterfactual supporting, but not so all generalizations. Nevertheless, we shall follow the Horgan/Tienson usage.

[3] It might be said that rules are abstract objects, hence cannot be involved in causal processes; only *representations of rules* can be involved in causal processes. I don't wish to make an issue of this, so I have no objections to anyone taking my statements about governance by rule to be elliptical for the statements about governance by representations of rules.

[4] In their 1989 paper it may have seemed that Horgan and Tienson assumed that Smolensky's tensor product theory provided a sufficient response to Fodor and Pylyshyn. In personal correspondence, Horgan has told me that they believe that much more remains to be done in order to meet Fodor and Pylyshyn's objection.

[5] Horgan and Tienson do not explain how we are to align semantic interpretations of input and output patterns of activation values with intentional concepts, such as belief and desire, so I propose to read these rule forms in as philosophically neutral a way as is possible. I suggest that they might be read as saying that, if the network is in input psychological state $RI_i(00)$, then it will go into output psychological state $RI_o(**)$. I think such a construal will suffice for present purposes.

[6] It might be suggested that we can use infinitely many activation values per node, either in the form of infinitely many integer valued activation values or in the form of real valued activations, to give rise to infinitely many representations. This is not a standard usage of connectionist activation values and a more detailed accounting of the proposal would be desirable. Further, Terry Horgan has said, in personal conversation, that he does not wish to propose the use of real-valued, or analogue, representations. In any case, there seems to be no reason to say that such networks would not be governed by some representation-level equations. Here is a sketch of the way it would go. Allow the input activation values to represent something like a confidence scale, so, using a two-input-node-two-output-node network, one would have an input pattern of activation, say, 35 92 represent a confidence of 35 92 that a dog is present (whatever that might mean exactly) and an input pattern of activation, say, –678 –456 represent a confidence

of −678 −456 that a dog is present (whatever that might mean exactly). This idea is related to the idea of computation by interactive activation (cf McClelland and Rumelhart 1981, Rumelhart and McClelland 1982). Exactly how (and why) these infinitely many representations might be parlayed into a psychological model does not matter for present purposes. In this sort of scheme, the behavior of the network will be governed by a representation-level rule such as,

If the confidence in the presence of a dog is a_1 a_2,
then the confidence that one should flee is a_3 a_4,

where the values of a_3 and a_4 are given as a function of the values of a_1 and a_2. I spare the reader a technical definition of the function relating a_1, a_2, a_3 and a_4. The function relating a_1, a_2, a_3 and a_4 may be a real-valued function, hence not tractable in the sense of not Turing computable, but it will nonetheless be quasi-exceptionless. The function would define an intractable, quasi-exceptionless rule governing the network. This is sufficient to undermine the central concern of RWR.

[7] Since Rogers' text was originally published in 1967, computer scientists have seen fit to drop the deterministic condition in the interests of the development of complexity theory and, more specifically, the theory of NP-completeness (cf. Garey and Johnson 1978).

[8] Cutland does not mention a finiteness requirement for Markov algorithms (Cutland 1980, pp. 64–5), but presumably this is only an infelicity arising from the very abbreviated discussion of this sort of computational device.

[9] Machtey and Young do not explicitly require that Turing machine programs contain only finitely many instructions, but this requirement follows from some of their other requirements. They assume that there are only finitely many tape symbols, that there are only finitely many state symbols for the read-write head, and that no two instructions can begin with the same state symbol and tape symbol. This entails that there can only be finitely many instructions in any Turing machine.

[10] Here one might wish to try to introduce a theory of analogue representations into connectionism. These representations would not be digital, hence not literally Turing-computable, hence computations performed over them would not be tractable in the sense of computable by Turing-equivalent device. As mentioned above, even if such a theory of representation were successfully developed for connectionist networks, there would still be intractable, quasi-exceptionless rules governing the networks.

[11] John Heil has something like these objections in mind when he has commented to me, "You seem to think that any finite task must count as "tractable" (on any plausible reading of 'tractable'). But consider a task that is finite but would take longer than the history of the universe to complete. Is such a task tractable? Is there any sense in talking of it as 'doable in principle'? I'm doubtful. I'm particularly doubtful if your aim is the modelling of human cognition" (Heil, personal correspondence).

REFERENCES

Cutland, N. J.: 1980, *Computability: An Introduction to Recursive Function Theory*, Cambridge University Press, Cambridge.

Davis, M.: 1982, *Computability and Unsolvability*, Dover, New York.

Dreyfus, H. and Dreyfus, S.: 1988, 'Making a Mind Versus Modeling the Brain: Artificial Intelligence Back at a Branchpoint', in S. R. Graubard (ed.), *The Artificial Intelligence Debate: False Starts, Real Foundations*, MIT Press, Cambridge, Massachusetts, pp. 15–43.

Fodor, J. and McLaughlin, B.: 1990, 'Connectionism and the Problem of Systematicity: Why Smolensky's Solution Doesn't Work', *Cognition* **35**, 183–204.

Fodor, J. and Pylyshyn, Z.: 1988, 'Connectionism and Cognitive Architecture', *Cognition* **28**, 3–71.

Garey, M. and Johnson, D.: 1979, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, New York.

Horgan, T. and Tienson, J.: 1988, 'Settling into a New Paradigm', *Connectionism and the Philosophy of Mind: Proceedings of the 1987 Spindel Conference, Southern Journal of Philosophy* **26**, 97–113, supplement.

Horgan, T. and Tienson, J.: 1989, 'Representations Without Rules', *Philosophical Topics* **17**, 147–174.

Horgan, T. and Tienson, J.: 1990, 'Soft Laws', *Midwest Studies in Philosophy: The Philosophy of the Human Sciences* **15**, 256–79.

Horgan, T. and Tienson, J.: 1992, 'Cognitive Systems as Dynamical Systems', *Topoi* **11**, 27–43.

Machtey, M. and Young, P.: 1978, *An Introduction to the General Theory of Algorithms*, North Holland, New York.

McClelland, J. L. and Rumelhart, D. E.: 1981, 'An Interactive Activation Model of Context Effects in Letter Perception: Part I. An Account of Basic Findings', *Psychological Review* **88**, 375–407.

McClelland, J. L., Rumelhart, D. E. and the PDP Research Group. (eds.): 1986, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vols. 1, 2, MIT Press, Cambridge, Massachusetts.

Rogers, H.: 1987, *Theory of Recursive Functions and Effective Computability*, MIT Press, Cambridge, Massachusetts.

Rumelhart, D. E. and McClelland, J. L.: 1982, 'An Interactive Activation Model of Context Effects in Letter Perception: Part 2. The Contextual Enhancement Effect and Some Test and Extensions of the Model', *Psychological Review* **89**, 60–94.

Smolensky, P.: 1989, 'On Variable Binding and the Representation of Symbolic Structures in Connectionist Systems', Technical Report CU-CS-355-87, Department of Computer Science, University of Colorado at Boulder, February.

Department of Philosophy
Central Michigan University
Mt. Pleasant, MI 48859
e-mail:35C4UB2@CMUVM.BITNET