WILLIAM RAMSEY AND STEPHEN STICH

CONNECTIONISM AND THREE LEVELS OF NATIVISM

ABSTRACT. Along with the increasing popularity of connectionist language models has come a number of provocative suggestions about the challenge these models present to Chomsky's arguments for nativism. The aim of this paper is to assess these claims. We begin by reconstructing Chomsky's "argument from the poverty of the stimulus" and arguing that it is best understood as three related arguments, with increasingly strong conclusions. Next, we provide a brief introduction to connectionism and give a quick survey of recent efforts to develop networks that model various aspects of human linguistic behavior. Finally, we explore the implications of this research for Chomsky's arguments. Our claim is that the relation between connectionism and Chomsky's views on innate knowledge is more complicated than many have assumed, and that even if these models enjoy considerable success the threat they pose for linguistic nativism is small.

1. INTRODUCTION

About 25 years ago, Noam Chomsky offered an argument aimed at showing that human beings must have a rich store of innate knowledge, because without such innate knowledge it would be impossible for children to learn a language on the basis of the data available to them. This "argument from the poverty of the stimulus" has had an enormous impact in linguistics, cognitive science, and philosophy. Jerry Fodor has described it as "the existence proof for the possibility of cognitive science . . . [and] quite possibly the only important result to date".[1] Hornstein and Lightfoot have urged that the argument serves as the foundation for most current work in linguistics.[2] And a number of authors, including Chomsky himself, have maintained that the argument from the poverty of the stimulus shows that empiricist theories of the mind are mistaken and that "the only substantive proposal to deal with the problem of acquisition of knowledge of language is the rationalist conception . . . ."[3]

During the last few years, however, a new research program, often called 'Connectionism' or 'Parallel Distributed Processing' (PDP), has attracted considerable attention in cognitive science. Connectionist

Synthese **82**: 177–205, 1990.

models of cognitive processes differ in many ways from earlier accounts commonly adopted by Chomskians. What makes them important for our purposes is that they employ powerful new learning techniques that enable systems to acquire complex and subtle skills in a wide variety of domains, without the assistance of large amounts of pre-programmed information. Very early on it was clear that the existence of these strikingly powerful learning strategies was a *prima facie* challenge to Chomsky's nativism. One observer, for example, comments that connectionism "sustains the vision of larger machines that are built on the same principles and that will learn whatever is learnable *with no innate disposition to acquire particular behaviors*" (italics ours).[4] If connectionist models invoking 'back propagation' or other learning algorithms can quickly acquire a large variety of complex skills without the help of 'innate' knowledge, it is natural to wonder whether they might not be able to acquire linguistic skills of the sort Chomsky argued could only be acquired by systems richly endowed with linguistic information at the outset. Motivated in part by just such anti-nativist suspicions, a number of investigators have begun to explore the possibility that connectionist models might acquire natural language syntax, phonology, semantics, and other features of linguistic ability. These efforts to build connectionist networks that learn aspects of natural language are very recent, and the results to date are both fragmentary and controversial. It is too early to venture a prediction on how successful they ultimately will be.

In this paper, our aim will be to explore the relation between connectionism and Chomsky's arguments for the existence of innate knowledge. Along the way, we propose to defend a pair of interrelated conclusions. The first is that there are actually three versions of Chomsky's poverty of the stimulus argument, which make increasingly strong claims about the nature of the cognitive endowments required for learning language. Though the three versions of the argument are often run together in the literature, it is essential to pull them apart if we are to be clear on the bearing that connectionist research might have on nativism. Our second conclusion is that the relation between connectionism and nativism is considerably more complex than many have assumed. There are various connectionist research programs which would, if successful, undermine all three versions of the Chomskian argument. However, the weakest version of the argument, whose conclusion is a doctrine that we will call *minimal nativism*, is easy to reconstruct in a way which will withstand any findings that may be forthcoming from connectionist research. A second version of the argument, aimed at establishing a stronger claim that we will call *anti-empiricism*, can also readily be reconstructed in the face of any foreseeable connectionist successes. However, both Chomsky's formulation of this argument and the reconstruction we will sketch require some sophisticated linguistic data. There has been a fair amount of linguistic research aimed at assembling the sort of data Chomsky's formulation of the argument requires. It is plausible to suppose that if the data needed in Chomsky's formulation are forthcoming, then linguists will be able to find an analogous body of data of the sort required by our reformulation. But, of course, there can be no guarantee on this point until the work is done. The third version of Chomsky's argument seeks to establish the strongest of the three nativist claims, the one we will call *rationalism*. Here there are indeed imaginable connectionist achievements that would show the conclusion of the argument to be false. However, there are also many ongoing connectionist explorations of language learning whose success would be fully compatible with rationalism. The bottom line, then, is that while connectionism challenges Chomskian nativism in a variety of ways, it may well turn out that even the strongest version of nativism is compatible with spectacular connectionist successes in the modeling of language acquisition.

The remainder of the paper will be organized as follows. In Section 2 we will set out the three versions of the poverty of the stimulus argument. In Section 3 we will offer an introductory overview of recent connectionist research and a quick survey of ongoing efforts to get connectionist devices to learn aspects of natural language. In Section 4 we will explore the ways in which the success of these efforts would bear upon the three versions of Chomsky's argument.

2. THREE VERSIONS OF THE POVERTY OF THE STIMULUS ARGUMENT AND THREE LEVELS OF NATIVISM[5]

What changes occur when a child learns a language? The answer, of course, is that there are many changes. The most conspicuous is that the child is able to understand the language, to communicate with it, and to use it for all sorts of purposes. There are also less obvious changes. Once a child has mastered a language, he is capable of making a wide range of judgments about the properties and relations of expressions in the language. Thus, for example, speakers of English are normally capable of judging whether any arbitrary sound sequence constitutes a grammatical sentence of English, and if it does, they are capable of judging whether or not it is ambiguous; they are also capable of judging whether two arbitrary sentences are related as active and passive, whether they are related as declarative and yes-no questions, whether one is a paraphrase of another, whether one entails the other, and so on for a number of additional linguistic properties and relations. These sorts of judgments, or 'linguistic intuitions' as they are more typically called, have played a central role in generative linguistics since its inception.

It is, Chomskians maintain, a perfectly astounding fact that ordinary speakers of a language can make a practically infinite number of judgments about the grammatical properties and relations of expressions in their language. The most plausible explanation of this ability, they urge, is that speakers have a generative grammar of their language – an explicit system of rules and definitions – stored somewhere in their mind or brain. On Chomsky's view, "the mature speaker has internalized a grammar with specific properties . . . [and] in understanding speech he makes use of this grammar to assign a precept to a signal".[6] "To know a language . . . is to be in a certain mental state . . . consisting of a system of rules and principles".[7] This system of internally represented rules guides the complex and prolific linguistic judgments that the speaker is capable of making. It is also used, in various ways, in the more ordinary processes of language production and comprehension. If there is no internally represented grammar, Chomsky and his followers urge, then it is something of a mystery how speakers are capable of having the linguistic intuitions they have. The mentally stored grammar that is posited is not, of course, accessible to consciousness. Speakers cannot tell us the rules of the grammar represented in their brains any more than they can tell us how they go about recognizing faces or recovering salient information from memory. But if speakers do have an internally represented grammar, then a natural goal for the generative grammarian would be (and has been) to discover that grammar – the grammar that is 'psychologically real'.

The argument for the thesis that speakers have an internally represented generative grammar of their language has the form of an inference to the best explanation:

I know of no other account that even attempts to deal with the fact that our judgments and behavior accord with and are in part explained by certain rule systems . . . .[8]

Later, we will explain why many connectionists believe their models call this thesis into question. For now, however, let us assume that Chomsky is right and that speakers do indeed have a mentally stored grammar of their language. We can then develop the three versions of the poverty of the stimulus argument against the background assumption that the mechanisms subserving language acquisition must be able to produce the grammar that the child comes to internally represent.

2.1. *The Argument for Minimal Nativism*

The weakest version of the poverty of the stimulus argument begins with the observation that, during the time span normally required to learn a language, a child is exposed to only a very impoverished sample of often misleading linguistic data. This 'poverty of the stimulus' is due to three important aspects of the 'primary linguistic data':

(1) The set of sentences that a competent speaker of a language can use, comprehend, and offer linguistic intuitions about is vastly larger than the idiosyncratic set of sentences to which children are exposed in the course of learning a language.

(2) While learning their language, the speech children hear does not consist exclusively of complete grammatical sentences. Rather, they are typically exposed to a large assortment of non-sentences, including slips of the tongue and incomplete thoughts, samples of foreign languages, and even intentional nonsense. Thus, the data the child has available for learning to tell sentences from non-sentences are remarkably messy.

(3) Children, unlike linguists, are rarely given any indication that certain queer and complex sentences are ungrammatical, that certain pairs of sentences are paraphrases of one another, and so on. Hence, many sorts of data that linguists rely upon heavily in deciding between competing grammars – such as data derived from speakers' linguistic intuitions – are not available to the child.

That children can acquire a grammar at all on the basis of this sort of data requires that they have a learning mechanism of some sort in place before the acquisition process begins. A video recorder exposed to the primary linguistic data that a child is exposed to does not end up with an internally represented grammar. Nor, for that matter, does a puppy or a young chimpanzee. The cognitive system which the child brings to the task of language learning must be able to go from a limited and messy sample of data to a grammar that generates most of the sentences in the data, and a huge number of additional sentences as well. And any cognitive system capable of projecting beyond the data in this way is going to be reasonably sophisticated. So, given our assumption that children do in fact end up with an internally represented grammar, the 'poverty of the stimulus' seems to require that children come to the language learning task with an innate learning mechanism of some sophistication. Moreover, despite exposure to significantly different samples of data, different children in the same linguistic community end up having essentially the same linguistic intuitions, and thus, it is plausible to suppose, essentially the same internalized grammar. Nor is there any evidence that children have any special predisposition to learn the language of their biological parents. Chinese children raised in an English-speaking environment learn English as easily as English children do. All of this suggests that the innate learning mechanisms that enable children to internalize the grammar of the language spoken around them are much the same in all children.

The crucial step in this first version of the poverty of the stimulus argument is the observation that if the child's innate learning mechanism is to accomplish its task, it must have a strong bias in favor of acquiring certain grammars and against acquiring others. This is because the data that the mechanism has been exposed to by the time grammar acquisition is complete is equally compatible with an indefinitely large class of grammars, many of which will depart in significant ways from the grammar that the child actually attains. The acquisition mechanism must project from the limited data it has available to a correct grammar – one that classifies sentences the way others in the linguistic community classify them. Thus it must somehow reject the indefinitely large class of *in*correct grammars that are equally compatible with the data. The thesis that we will call *minimal nativism* is simply the claim that the child approaches the task of language acquisition with an innate learning mechanism that is strongly biased in favor of certain grammars and against others. But, of course, to say that the innate learning mechanism is biased in favor of certain grammars and against others does not commit us to any particular account of the mechanism underlying this bias. It is on just this point that the three levels of nativism differ.

Minimal nativism merely insists that the bias must be there. The higher levels of nativism make increasingly strong claims about the mechanism responsible for the bias.[9]

Before moving on to the next version of the argument, it is important to make clear exactly what does and does not follow from minimal nativism. One might think that by establishing the existence of a strongly biased innate learning mechanism, Chomsky has succeeded in undermining the empiricist conception of the mind. But this would be a mistake. For even the staunchest empiricist would readily agree that learning requires sophisticated innate mechanisms and biases. As Quine reminds us, the empiricist "is knowingly and cheerfully up to his neck in innate mechanisms of learning readiness".[10] If Chomsky's argument is supposed to undermine empiricism, then it must say something about the nature of these mechanisms and biases which calls into doubt the empiricist conception of the mind.

2.2. *The Argument Against Empiricism*

At first blush, it might be thought that it would be impossible to argue against *all* empiricist accounts of the mind. For while Chomsky might show that on one or another specific empiricist theory, the mind could not reliably produce the right grammar on the basis of the primary linguistic data, it would always be open to the resourceful empiricist to construct another theory, still adhering to empiricist principles, though diverging in one way or another from the particular empiricist theory that has been refuted. However, there is in Chomsky's writings an ingenious idea for circumventing this problem and refuting all empiricist theories in one fell swoop. We'll call this idea 'the Competent Scientist Gambit'. The basic idea is to portray a learning mechanism that is at least as powerful as anything dreamt of in the empiricist conception of the mind, and then argue that such a learning mechanism could not do what the child does. If this can be shown, then all empiricist theories will fall together. The 'learning mechanism' Chomsky suggests is a competent, rational scientist.

Suppose that we were to pose for such a scientist the task at which the child's mind is so adept. We will give the scientist a typical set of primary linguistic data drawn from some actual human language. Her job will be to discover the grammar of that language – the grammar that children exposed to those data will come to internally represent.

In going about the business of constructing and testing hypotheses about the grammar she is trying to discover, the scientist will be able to exploit any inferential strategy that would be permitted by any account of the mind compatible with empiricist strictures. She can record data, do sophisticated data analysis, think up imaginative hypotheses (or mundane ones) and test those hypotheses against the data available to her. Moreover, it is open to her to employ the sorts of methodological principles and intuitions typically employed in empirical theory construction and selection. In discussions of those methodological considerations, simplicity often looms large, and from time to time we will use the term 'simplicity' as a convenient label for the whole package of methodological principles and intuitions that a competent scientist has available.

There is, however, one thing that the competent scientist is not allowed to do. She is not allowed to learn the language from which the primary linguistic data are drawn. There is, of course, no reason to think that the scientist could not learn the language on the basis of that data. She is a normal human, and we are providing her with just the sort of data that generally suffices for normal humans to learn a language. The point of the prohibition is simply that if she were to learn the language, she would then have access to data that the child does not have. She would have her acquired linguistic intuitions about the grammaticality of sentences not presented in the data, as well as her intuitions about ambiguities, about paraphrases, and so on. But if her challenge is to try to do what the child does, then it is obviously unfair for her to use information not available to the child. Clearly it is absurd to suppose that in order to learn his language the child must first learn it, and then generate the data necessary for him to learn it.

We are supposing that after exposure to a decade or so of primary linguistic data from any natural language, the child succeeds in constructing a grammar that projects well beyond his data, and does so correctly, where the standard of correctness is set by the senior members of the child's linguistic community. If the scientist is to match the child's feat, she too must make a monumental projection from the data available to her, and come up with the grammar that has been internalized by those who are producing the data. Chomsky's contention is that given only the information embodied in the primary linguistic data, along with the methodological resources available to her, the competent scientist could not reliably do what the child does. That is, the scientist could not discover the grammar the child comes to internally represent when learning a language.

It is important to understand exactly what is being claimed when Chomsky makes this assertion. Chomsky does not deny that the competent scientist could *think up* the right grammar. Of course she could. *Ex hypothesis* she is intelligent, creative, and resourceful, so if she couldn't think up the right grammar, no one could. However, there is a sense in which this very intelligence and creativity is the scientist's undoing. For just as there is every reason to believe she can think up the *right* grammar – the one the child actually ends up with – so too there is every reason to believe she can think up an endless variety of *wrong* grammars that do not project from the data in the way the child's grammar does. The crucial contention for this version of the poverty of the stimulus argument is that *the methodological resources a scientist has available will not suffice to motivate the proper selection*. Even with the use of criteria such as simplicity, the scientist would still be plagued by an embarrassment of riches. In saying that the scientist would be incapable of 'coming up with' the right grammar, what is meant is that the scientist will have no reliable way of locating the right grammar in the space of possible grammars that are compatible with the limited data she has available.[11]

It now should be clear how the Competent Scientist Gambit is intended to undermine the empiricist conception of learning. It is plausible to view the competent scientist as a strong and generous characterization of the empiricist mind. (Indeed, there will be many things a competent scientist can do that the sort of mind conjured by the Classical Empiricists cannot.) Hence, if the competent scientist is not up to the task, then no learning mechanism compatible with empiricist principles will be adequate for the task of language acquisition. If it can be shown that something at least as resourceful as the empiricist mind would fail at language learning, Chomsky will have succeeded in showing that the empiricist conception of the mind must be mistaken.

Of course for all of this to work, some additional argument is going to be needed. What needs to be shown is that the set of methodological principles and biases available to a competent scientist will not be adequate for successful projection from the primary linguistic data to the grammar of the language from which the data are drawn. One way to show this would be to produce a pair of grammars with the following features:

(i) on all intuitive measures of simplicity the grammars are comparable;

(ii) the grammars make essentially the same judgments about linguistic phenomena that are likely to show up in the primary linguistic data; and

(iii) the grammars make significantly different judgments about linguistic phenomena that are not likely to show up in the primary linguistic data.

If there are examples of this sort, our competent scientist will be unable to choose between the grammars. Since the grammars are both compatible with any plausible body of primary linguistic data, she cannot use the data to rule one out. And since they are both comparably simple, methodological considerations will be of no help. If, in these cases, language learners regularly project in the right way, it follows that the mechanisms responsible for language learning must be more powerful than the empiricist conception of the mind will allow.

In recent years, there has been a fair amount of work in linguistics aimed at compiling examples of just this sort. For example, Hornstein and Lightfoot[12] sketch a case in which the choice between two very different, though comparably simple grammars turns on the paraphrase relations among sentences like (1)–(3):

(1) She told me three funny stories, but I didn't like the one about Max.

(2) She told me three funny stories, but I didn't like the story about Max.

(3) She told me three funny stories, but I didn't like the funny story about Max.

On one of the grammars under consideration, (2) would be considered a paraphrase of (1), though (3) would not. The other grammar correctly entails that both (2) and (3) might be paraphrases of (1). It is, Hornstein and Lightfoot maintain, very unlikely that every child who successfully learns English will have been exposed to primary linguistic data containing evidence about these sorts of relatively abstruse facts concerning paraphrase. If this is right, and if the only sorts of evidence that would suffice to distinguish between the two grammars are comparably abstruse, then our competent scientist is in trouble. Since she is intelligent and resourceful, she will be able to think up both grammars. Since neither grammar is simpler nor superior on other methodological grounds, such considerations will not assist her in making the correct choice. And, unlike the real linguists who actually did worry about the choice between these two grammars, she does not have, and cannot get, the kind of data that would enable her to make the right choice.

The argument just sketched is, of course, very much hostage to the linguistic facts. For the argument to be persuasive there must be a substantial number of examples in which the choice between two equally simple and natural grammars can be made only by appealing to the sort of abstruse evidence that is unlikely to be found in the primary linguistic data. There is by now a substantial collection of plausible cases in the literature.[13] If these cases survive critical scrutiny, Chomsky and his followers will have gone a long way toward making their case against empiricism.

This brings us to the conclusion of the second version of the poverty of the stimulus argument, a doctrine we shall call *anti-empiricism*. This doctrine maintains not only that the innate language learning mechanism must have strong biases, but also that these biases are not compatible with the account of mental mechanisms suggested by even a very generous characterization of the empiricist mind. Anti-empiricism makes a negative claim about the language learning mechanism – a claim about what its biases are not. The third version of the poverty of stimulus argument aims at establishing a positive claim about the way the language learning mechanism does its job.

2.3. *The Argument for Rationalism*

If the empiricist conception of the mind cannot account for the facts of language learning, what sorts of accounts of the mind can? One way of approaching this question is to focus on exactly why it was that our hypothetical scientist could not do what the child does. The problem was not that she could not think up the right grammar, but rather that she could also think up lots of wrong grammars that were equally simple and equally compatible with the data, and she had no way to decide among them. Confronted with this problem, one strategy that might enable the scientist to duplicate the child's accomplishment would be to narrow the range of grammars she must consider. Suppose it were the case that all the correct grammars of human languages – all the ones that speakers actually have represented in their heads – shared certain properties. If this were so, then the scientist's work would be greatly facilitated if she were informed about these properties at the outset. For then she would never have to consider any of the grammars that do not share the 'universal' features of all human grammars. The richer the collection of universal features, the stronger the constraints they will impose on the class of grammars that the scientist need consider; and the stronger the constraints, the easier her task will become.[14] What does all this suggest about the child's mind? The obvious hypothesis to extract from the analogy between the child's task and the scientist's is that the child's mind comes equipped with information about linguistic universals – biases that are applicable only in the area of language acquisition – that enable it to pick out the right grammar by narrowing the search space. On this hypothesis, the child begins with a rich body of innate information about language which serves to define the class of all human languages. The relatively impoverished environmental stimulus is "viewed as only a trigger; much of the ability eventually attained is determined by genetically encoded principles, which are triggered or activated by environmental stimulus rather than formed by it more or less directly".[15] Clearly this hypothesis goes well beyond the thesis that the biases built into the innate language learning mechanism are non-empiricist. As John Searle notes, "Chomsky is arguing not simply that the child must have 'learning readiness', 'biases', and 'dispositions', but that he must have a *specific* set of linguistic mechanisms at work."[16] Moreover, this domain specificity of innate mechanisms has been a traditional feature of rationalist conceptions of the mind. For Chomsky and his followers, the central argument for the claim that the child has domain specific language learning biases is, once again, an inference to the best explanation – it is "the only substantive proposal to deal with the problem of acquisition of knowledge of language".[17] And prior to the emergence of connectionism, Chomsky's argument was surely very plausible. Once we realize the difficulties facing the child, it is no easy matter to imagine how he could possibly solve the projection problem and end up with the right grammar, unless he approached the task with a rich set of constraints specifically tailored to the task at hand. The thesis that the innate language learning mechanism embodies such constraints is the conclusion to be drawn from the third version of the poverty of the stimulus argument. We'll call this view *rationalism*.

We've now completed our reconstruction of the three versions of the poverty of the stimulus argument and the conclusions that have been drawn from them. In Section 4 we will explore the ways in which connectionism might be thought to challenge these arguments. Before getting to that, however, we'll need to give a quick sketch of connectionism, and review some recent attempts to study linguistic phenomena in a connectionist framework.

3. AN OVERVIEW OF CONNECTIONIST RESEARCH ON LANGUAGE

Connectionism is a new style of cognitive modeling that has emerged during the last decade. Connectionist models consist of networks built from large numbers of extremely simple interacting units. Inspired by neuronal architecture, connectionist units are typically linked in such a way that they can excite or inhibit one another by sending activation signals down interconnecting pathways. Networks commonly involve a layer of input units, a layer of output units, and one or more intermediate (or 'hidden') layers, linked by weighted connections through which a wave of activation travels. When the processing proceeds in only one direction, as is the case with 'feed-forward' networks, units modify and transfer the activation signal only to subsequent units and layers. In other, more complicated networks, activation may involve feedback loops and bi-directional communication between nodes, comprising what are often referred to as 'recurrent' networks. The units themselves may have threshold values, which their total input must exceed for activation. Alternatively, they may act in analog fashion, taking an activation value anywhere between 0 and 100%. Connecting links have varying weights or strengths, and the exact nature of the activation signal transferred from one unit to another (that is, its strength and excitatory or inhibitory value) is typically a function of the connection weight and the activation level of the sending unit.

This architecture supports a style of computation quite unlike that exploited by earlier cognitive models. For the most part, pre-connectionist model builders have presupposed computational architectures that perform operations best described as 'symbol manipulations'. In such systems, information is generally stored in distinct locations separate from the structures performing computational operations. Information processing in such devices consists of the manipulation of discrete tokens or symbols, which are relocated, copied, and shuffled about, typically in accordance with rules or commands which are themselves encoded in a manner readily discernible by the system.

Connectionist information processing diverges from these earlier models in many ways. Perhaps the most striking aspect of connectionist information processing is that it typically does not involve anything like the manipulation of distinct symbolic tokens. While connectionist modelers sometimes invoke notions of representations to characterize elements of their networks, connectionist representations are generally not at all like the discrete symbolic entities found in classical architectures. This is especially true when the model employs 'distributed representations', where the same set of individual units and weights are used to encode divergent bits of information.[18] Another notable difference between connectionist models and earlier cognitive models is that in connectionist models the distinction between structures that store information and structures that process information is virtually non-existent. Information is 'stored' in the connection weights between individual units, which serve as central elements in the processing as well. Hence, familiar notions of stored programs or autonomous command structures which govern computational processing seem to have no place in connectionist architecture.

These differences loom large in the debate over the psychological reality of linguistic rules. As we saw in Section 2.1, Chomsky's formulations of the poverty of the stimulus arguments presupposes that when a child has learned a language he or she ends up with an internally represented generative grammar – typically a set of re-write or production rules each of which consists of a sequence of distinct symbols. Pre-connectionist cognitive models, which view cognition as symbol manipulation, are entirely comfortable with this view. But connectionist models, particularly those exploiting highly distributed representations and non-modular computational strategies, cannot readily accommodate the sorts of symbolic rules posited by generative grammarians. In defense of the claim that linguistic abilities are subserved by an internally represented grammar, Chomsky offered an inference to the best explanation argument. Appeal to internalized grammatical rules was not only the best way to explain linguistic judgments and behavior, Chomsky maintained, it was the *only* explicit, well-developed hypothesis that had ever been suggested. Prior to the emergence of connectionism, that argument had considerable plausibility.[19] If, however, it turns out that connectionist models can account for much the same range of data about linguistic intuitions and linguistic behavior, it will no longer be possible for Chomsky and his followers to claim that their internalized rule explanations are the "only game in town".

Since connectionist information processing is governed by connection weights between units, the computations can be altered simply by changing the value of these weights. Connectionist researchers realized early on that if weight changes could be executed in a purposeful manner, then these models would manifest a form of learning that seems biologically plausible, and quite revolutionary from a computational perspective. Recent developments have overcome past difficulties in multi-layer weight adjustment, and there are now very powerful learning strategies that enable connectionist networks to, in a sense, program themselves. Perhaps the most widely used learning algorithm is the 'generalized delta rule' or 'back propagation', developed by Rumelhart, Hinton, and Williams.[20] On this learning strategy, a network undergoes a training period during which it is presented with a series of inputs and allowed to produce an output for each presentation. A comparison is made between the actual output and a target output for each presentation, resulting in an error signal. This signal is subsequently propagated back through the network, adjusting weights in accordance with the learning algorithm. Because the weights are fixed after training, the system is subsequently able to make 'educated' responses to new inputs that were not presented during the learning period. The success of most models is determined by how well they perform such generalizations within a particular task domain.[21]

So much for our general overview. There are many other styles of connectionist processing and learning, but this should suffice to give a sense of the basic elements of the new paradigm. Let's turn now to the growing body of connectionist research devoted to developing models of language processing and language acquisition. Much of this research has been motivated by increasing skepticism about Chomsky's account of language acquisition, and by the suspicion that language processing and acquisition might be more naturally explained by models with connectionist architectures. Prior to the emergence of connectionism, Chomsky often stressed that "it is difficult to imagine how the vague suggestions about conditioning and associative nets that one finds in philosophical and psychological speculations of an empiricist cast might be refined or elaborated so as to provide for attested competence . . . ."[22] Many connectionists believe that their new computational tools overcome such failures of the imagination, and have developed impressive models aimed at making the point.

A typical model of this sort is PARSNIP, developed by Hanson and Kegl. [23] This is an auto-associator network[24] that was trained on three sets of syntactically tagged natural language sentences. Beginning with the assumption "that natural language reveals to the hearer a rich set of linguistic constraints . . . that serve to delimit the possible grammars that can be learned" (p. 108), the modelers found that a network trained to produce veridical copies of input could also "induce grammar-like behavior" while performing various linguistic tasks.[25]

The network learned to produce correct syntactic category labels corresponding to each position of the sentence originally presented to it, and it was able to generalize to another 1000 sentences which were distinct from all three training samples. PARSNIP does sentence completion on sentences, and also recognizes novel sentence patterns absent from the presented corpus. One interesting parallel between PARSNIP and human language users is the fact that PARSNIP correctly reproduces test sentences reflecting deep center-embedded patterns which it has never seen before while failing to reproduce multiply center-embedded patterns.[26]

While Hanson and Kegl concede that their model has certain psychologically implausible features (such as insensitivity to temporal factors), they maintain that

there are important parallels between the task given to PARSNIP and the task that arises for children as they learn a natural language. Both PARSNIP and the child are only exposed to sentences from natural language, they both must induce general rules and larger constituents from just the regularities to which they are exposed, both on the basis of only positive evidence. PARSNIP's ability to generalize knowledge of constituent structure has been extracted from its experience with natural language sentences.[27]

A number of connectionist models attempt to account for aspects of language that have been difficult to capture in more conventional rule-based systems. It appears that sensitivity to several different sources of information (such as cues from phonetic, semantic, and contextual factors) is much easier to implement in connectionist networks with distributed encodings and parallel processing. One system exploiting this type of architecture was designed by McClelland and Kawamoto (1986) to assign correct case roles to constituents of sentences. The model invokes word order and semantic constraints to determine case assignments and to select contextually appropriate readings of ambiguous words. A similar but more complex model of semantic processing developed by St. John and McClelland (1988) learns mappings between words in particular contexts and concepts, and predicts additional meanings implicit in the sentence.

While these systems focus primarily on semantic aspects of language comprehension, a number of connectionist models have been developed to account for syntactic, phonological, and other non-semantic components of language processing. For example, Fanty (1985) has developed a connectionist parser that incorporates all levels of the parse tree at the same time, producing the surface structure of the sentence as its output. Other efforts at connectionist parsing include models by Cottrell (1985), Waltz and Pollack (1985), Selman and Hirst (1985), and Charniak and Santos (1986). Rumelhart and McClelland (1986b) have produced a network designed to model the acquisition of English past-tense verbs. The most intriguing feature of this model is its ability to replicate putative aspects of human past tense learning such as overgeneralization of regular past-tense forms to irregular forms without incorporating the sort of discrete symbolic rules commonly assumed to account for such phenomena.[28] Elman (1988) has produced a model that learns to divide an unbroken stream of input into phonemes, morphemes, and words, a capacity often claimed to be largely innate. The model also produced representations of lexical classes through exposure to word order alone, distinguishing nouns and verbs, for example, and arranging their representations into various semantic hierarchies.

It should be clear from this (by no means exhaustive) survey that connectionist language modeling is a robust and thriving area of research. As we noted at the outset, it is too soon to tell just how successful such work will ultimately be. However, our concern here is not to debate the superiority of connectionist models but to explore how the arguments for nativism will fare if connectionist models prove to be empirically accurate accounts of the mechanisms underlying language acquisition and linguistic competence. That is the issue we'll tackle in the section to follow.

4. CONNECTIONISM AND NATIVISM

In Section 2 we detailed three versions of the poverty of the stimulus argument that yield three distinct conclusions making progressively stronger nativist claims. There are two ways in which it might be thought that advances in connectionist language modeling could threaten those arguments. The first focuses on the output of the language acquisition process, the second on the nature of the process itself. We'll begin by sketching both of these challenges, and then go on to ask how much damage they do to each version of the argument from the poverty of the stimulus.

4.1. *The First Connectionist Challenge: Adult Competence Is Not Subserved by a Grammar*

As set out in Section 2, all three versions of the argument from the poverty of the stimulus begin with the assumption that when a person learns a language he or she ends up with an internally represented grammar of that language, where a grammar is taken to be a system of generative rules built out of an appropriate symbolic vocabulary. The Chomskian defense of this assumption is that "it's the only game in town" for explaining language competence. But, as we saw in Section 3, connectionist models don't readily accommodate the sorts of symbolic rules exploited by generative grammarians. Thus, if it turns out that connectionist models of adult linguistic competence can account for a wide range of linguistic judgments and abilities, Chomskians will no longer be able to claim that a theory positing an internalized grammar is the only option available. And if connectionist models of linguistic competence prove to be empirically *superior* to models invoking internalized grammars, the poverty of the stimulus arguments will have to do without the assumption that the output of the acquisition process includes an internally represented grammar.

4.2. *The Second Connectionist Challenge: Connectionist Learning Algorithms Can Model Language Acquisition*

All three versions of the poverty of the stimulus argument conclude that the mechanism responsible for language acquisition must be biased in favor of certain outcomes and against others. On the anti-empiricist version of the argument, the biases are claimed to be incompatible with the account of the mind envisioned in the empiricist tradition. On the rationalist version, the biases are further claimed to be specific to language and applicable only in the domain of language acquisition. But suppose it could be shown that a system using back propagation or another connectionist learning algorithm can do a good job at modeling some impressive part of the child's accomplishment in learning a language. Suppose, for example, that a connectionist acquisition model could mimic the language learner's projection from primary linguistic data to judgments about sentences that he or she has never heard. We might imagine the hypothetical connectionist acquisition model behaving as follows: When provided with a sample of primary linguistic data from any natural language (i.e., a large set of utterances of the sort that a child learning the language might be exposed to, perhaps accompanied by some information about the setting in which the utterance occurs) the model learns to distinguish grammatical sentences in that language from ungrammatical ones with much the same accuracy that a human learner does.

It might well be thought that the existence of such a model would refute all three versions of nativism. For, it might be argued, back propagation and other connectionist learning algorithms, far from being restricted to language, appear to be enormously general in their domain of application. Back propagation has been used successfully in training networks to perform very diverse tasks – from transforming written text into phonemes to distinguishing sonar echoes of rocks from those of undersea mines.[29] Thus the learning model we have imagined appears to pose a direct challenge to the doctrine we have been calling 'rationalism'. Moreover, connectionist learning algorithms like back propagation seem to be very much in the spirit of the simple, general-purpose learning mechanisms envisioned in the empiricist tradition. Historically, back propagation can be viewed as a variant on a simple learning rule suggested by Hebb.[30] And the (unmodified) 'delta rule' was first proposed by Sutton and Barto as part of their theory of classical conditioning.[31] So if the sort of connectionist acquisition model we have been imagining could actually be built, it would appear to pose a challenge to the doctrine we have been calling 'anti-empiricism'. It might even be urged that the existence of such a model would threaten minimal nativism, since back propagation and other connectionist learning algorithms seem remarkably free from biases of any sort. This may be what Sampson has in mind when he writes: "[T]he knowledge eventually stored in the system, in the pattern of weights, is derived entirely from the input". "The system's only contribution is to react in a passive, mechanical way to individual data items".[32]

These challenges make it sound like connectionism is on a collision course with Chomsky's nativism. On the one hand, if empirically successful connectionist models of adult linguistic competence can be built, a central assumption of the arguments from the poverty of the stimulus will be undermined. On the other hand, if connectionist learning algorithms can project from the primary linguistic data in the way the child does, the conclusions of all three arguments are threatened. However, on our view, even if things turn out well for connectionism, the challenge it will pose to Chomskian nativism will be far from devastating. It is true that in the wake of the connectionist achievements we have been imagining all three versions of the argument from the poverty of the stimulus would come unglued. But this alone would not refute any of Chomsky's nativist conclusions. As we'll see in the section to follow, we can readily formulate a new version of the argument for minimal nativism that sidesteps both connectionist challenges. The argument for anti-empiricism can also be reconstructed, as we'll see in Section 4.4, though it will require a sort of empirical evidence rather different from that exploited in Section 2.2. And, as we shall argue in Section 4.5, even Chomskian rationalism may turn out to be compatible with our hypothesized connectionist achievements.

4.3. *Connectionism and Minimal Nativism*

For argument's sake, let's grant that, despite Chomsky's argument to the contrary, the mechanisms subserving the linguistic skills of a competent speaker do not exploit an internally represented grammar. Rather, we'll suppose that a trained up connectionist network underlies a speaker's ability to judge sentences as grammatical or ungrammatical, etc. On this assumption, the job of the language acquisition mechanism will be to produce an appropriate network, one which judges sentences the way other speakers of the language do. The input available to the acquisition mechanism will be a typically untidy body of primary linguistic data drawn (mostly) from the language being acquired. And, of course, the network that is the output of the acquisition mechanism will have to respond appropriately to a vast class of sentences that the acquisition mechanism was never exposed to.

But now just as there are indefinitely many grammars which are comparably compatible with any given body of primarily linguistic data, though they diverge in the judgments they make about sentences not in that body of data, so too there are indefinitely many connectionist networks that agree, near enough, in their judgments about a given body of primary linguistic data, while diverging in their judgments about sentences not included in the data. Thus the language acquisition mechanism must somehow reject an indefinitely large class of networks all of which are comparably compatible with the data. To do this, obviously the mechanism will have to be strongly biased in favor of acquiring certain networks and against acquiring others. And that is just what minimal nativism maintains. All of this is quite independent of any assumption we might make about the algorithm used by the acquisition mechanism. If a connectionist acquisition mechanism using back propagation can in fact produce a trained up network that makes the right judgments about vast numbers of sentences not included in the primary linguistic data, then the conclusion to be drawn is not that minimal nativism is false, but rather that the learning algorithm being used is strongly biased in favor of certain projections and against others. This should be no surprise. The task of the language acquisition mechanism is an inductive learning task. And as Goodman and others demonstrated long ago, any successful inductive learning strategy must be strongly biased.[33]

4.4. *Connectionism and Anti-empiricism*

While minimal nativism claims merely that the language learning mechanism must be biased, the Chomskian argument for anti-empiricism maintains that simplicity and other methodological principles of the sort that a scientist might use in deciding among theories will not suffice in explaining the child's success in learning language. Recall that to make this point, the anti-empiricist argument outlined in Section 2.2 needed some sophisticated linguistic evidence. It required us to find cases in which a pair of grammars that are near enough equal with respect to simplicity and other methodological virtues also agree in their judgments about typical bodies of primary linguistic data. If these grammars disagree in their judgments about cases not likely to be found in the primary linguistic data, then the competent scientist trying to duplicate the child's accomplishment would have no way of deciding among them.

This argument for anti-empiricism clearly requires that the mechanism underlying linguistic competence be a grammar, since it rests upon very specific claims about the formal properties of grammars. But as we saw in Section 4.1, the success of connectionism would challenge this assumption. Since connectionist models of competence do not use anything like a grammar, the fact that different *grammars* are compatible with the data and equally simple would not suffice to establish anti-empiricism, if those connectionist models turn out to be right. Hence, the sort of connectionist models of linguistic competence that we have been imagining undermine the standard Chomskian formulation of the argument for anti-empiricism. This hardly constitutes a refutation of anti-empiricism, however, since it is possible to reconstruct an anti-empiricist argument parallel to Chomsky's which assumes that linguistic competence is subserved by a connectionist network.

Since we are assuming that adult linguistic competence is subserved by a connectionist network rather than a grammar, we will have to assemble cases in which a pair of connectionist networks have the following properties:

(i) the networks make much the same judgments about sentences likely to show up in the primary linguistic data;
(ii) the networks make significantly different judgments about sentences that are not likely to show up in the primary linguistic data; and
(iii) on intuitive measures of simplicity (and on other methodological grounds) the networks are much the same.

Since connectionist studies of language are of very recent vintage, and since many researchers in the area are skeptical about nativism, there has been no systematic effort to find such examples. Thus the data needed to secure our reconstructed anti-nativist argument are not available. But there is certainly no *a priori* reason to suppose that the evidence required cannot be found. And in assessing the threat connectionism poses for anti-empiricism, this last point is the crucial one. What it shows is that even if the suppositions in Sections 4.1 and 4.2 are correct, the truth of anti-empiricism will remain an open issue, to be decided by further empirical work. If the appropriate linguistic evidence *can* be found, and if the language acquisition mechanism is indeed a connectionist device exploiting back propagation, then the conclusion to be drawn is not that anti-empiricism is mistaken, but that the connectionist acquisition mechanism embodies biases different from those invoked in the empiricist tradition. More specifically, if the data turn out right, then the connectionist acquisition mechanism must be using something different from simplicity and other intuitive methodological principles. For *ex hypothesis* the acquisition mechanism is preferring one network to another, even though they are comparably simple and equally compatible with the data. Of course, if the data turn out the other way – if the appropriate linguistic examples are not to be found – then we will have no reason to regard anti-empiricism as true.

Before leaving the topic of anti-empiricism, there is one final point that needs attention. As we noted in Section 4.2, back propagation, the most widely used connectionist learning algorithm, was inspired by Hebbian learning rules and by work on classical conditioning. And while back propagation is significantly more sophisticated than Hebb's rule, or the (unmodified) delta rule invoked in the explanation of classical conditioning, it clearly shares a strong family resemblance with them. But, it might be argued, Hebb's rule, and the processes of classical conditioning are surely of a piece with the sort of mental processes that have been posited in the empiricist tradition. So if, as we have been assuming, a connectionist language acquisition device using back propagation could project from the data the way a child does, why should we not conclude that an empiricist acquisition device could succeed in learning language?

As we see it, the issue that is being raised here is how the notion of an 'empiricist' learning mechanism is best understood. Chomsky and his followers have adopted the competent scientist gambit as the acid test for empiricism. Any acquisition mechanism that can reliably do things a competent scientist cannot do does not count as an empiricist mechanism. And on this test it may well turn out that connectionist devices exploiting back propagation are not empiricist mechanisms. The alternative account of the notion of an 'empiricist' learning mechanism rejects the competent scientist standard, with its appeal to intuitive simplicity and other intuitive methodological considerations, and opts instead for the family resemblance criterion. On this account connectionist devices exploiting back propagation probably are empiricist mechanisms. As we see it, the dispute here is largely a verbal one. It will be an interesting and important fact if the competent scientist account of empiricism and the family resemblance account turn out not to coincide. But if this happens, who gets to keep the word 'empiricist' is a matter of very little moment.

4.5. *Connectionism and Rationalism*

Rationalism, as we have been using the term, is the thesis that the innate language learning mechanism embodies biases or constraints that are specific to the task of language learning, and of no use in other domains. The Chomskian justification for this thesis relies on the claim that there are no plausible alternatives. Thus in 1980, before the flourishing of connectionism, Wexler and Culicover wrote:

> At the present the constraints we need are quite specifically linguistic. More general theories would be intriguing, as insightful generalization always is, but until we have reason to believe the generalizations (or to formulate them coherently), we must remain skeptical.[34]

Here again, the connectionist achievements we've posited undermine the Chomskian argument. For, as we have noted, connectionist learning algorithms are anything but specifically linguistic. They have been used successfully in a wide variety of domains. So if a connectionist acquisition device could project from the primary linguistic data in the way the child does, Chomskians can no longer claim that rationalist acquisition models are the only game in town.

Undermining Chomsky's version of the argument for rationalism does not, however, show that rationalism is false; nor does it show that connectionism is incompatible with rationalism. For there are a great variety of connectionist learning devices that exploit back propagation. Some of them require idiosyncratic architectures or a great deal of pre-wiring and pre-tuning before they will do an acceptable job of learning in the task domain for which they are designed.[35] And as McClelland and Rumelhart note, such models are "clearly consistent with a rabidly nativist world view".[36] While connectionist research has produced learning strategies that are not domain specific, the extent to which these strategies can succeed in language acquisition without exploiting special architectures is currently unknown. If the only successful connectionist language acquisition devices are of a sort that require language specific architectures and/or language specific pretuning, then even the rationalist version of nativism will have nothing to fear from connectionism. Recently Rumelhart and others have been exploring ways in which connectionist learning algorithms themselves can be modified so as to bias learning in one direction or another.[37] If the best connectionist models of language acquisition exploit a learning algorithm that is particularly adept at language learning and largely useless in other domains, then again rationalism and connectionism will turn out to be comfortably compatible.

Of course, it is also conceivable that connectionist learning models will be able to duplicate significant aspects of the language learner's accomplishment without invoking idiosyncratic architectures, specialized pre-tuning or domain specific learning algorithms, and that much the same models will be able to master significant cognitive tasks in domains far removed from language. If such non-domain-specific models were to be developed, they would pose a genuine challenge for Chomskian rationalism.

5. CONCLUSION

The central claim of this last section has been that the putative incompatibility between connectionism and nativism has been much exaggerated. If adult linguistic competence is subserved by a connectionist network, and connectionist learning devices can duplicate the child's projection from primary linguistic data, all three versions of Chomsky's argument from the poverty of the stimulus will be undermined. However, parallel arguments for minimal nativism and anti-empiricism are easy to reconstruct. On our view, the argument for minimal nativism is entirely conclusive. The argument for anti-empiricism depends on empirical premises whose plausibility requires further investigation. There is no comparable reconstruction of the Chomskian argument for rationalism. However, if the only connectionist language acquisition models capable of projecting the way the child projects invoke language specific algorithms or architectures, then even rationalism will be sustained.

One final point is worth stressing. If it should turn out that non-domain specific models, like those envisioned at the end of Section 4.5, are capable of duplicating significant aspects of the child's accomplishment, and if the argument against empiricism can be successfully reconstructed, then our account of language acquisition would be located in the seldom explored terrain between rationalism and empiricism. It is here, perhaps, that connectionism may hold the most exciting potential for contributing to the nativism debate.

NOTES

[1] Fodor (1981), p. 258
[2] Hornstein and Lightfoot (1981b).
[3] Chomsky (1972), p. 88.
[4] Papert (1987), p. 8.
[5] Parts of this section are borrowed from Stich (forthcoming).
[6] Chomsky (1969), pp. 155–56.
[7] Chomsky (1980a), p. 48.
[8] Chomsky (1980b), p. 12.
[9] It is important to note that the relation between the primary linguistic data and a set of possible grammars is, in many ways, analogous to the abductive relation between evidential data and a set of different explanatory hypotheses. It is a truism in the philosophy of science that abductive inference – the projection from a body of data to an hypothesis that goes beyond the data – cannot be based upon the evidence alone. It requires an appeal to inferential principles or methodological criteria not included in the data. Similarly, since a child's primary linguistic data is compatible with a number of different grammars, his projection must be guided by some antecedent bias or set of constraints. For more on projection and language acquisition, see Gold (1967), Peters (1972), Wexler and Culicover (1980), and Morgan (1986).
[10] (1969, p. 95).
[11] Actually, this understates the difficulty that the scientist confronts since, as noted earlier, the primary linguistic data will typically be messy data, containing all sorts of sentences and sentence fragments that the correct grammar will not generate. So the task the scientist confronts is to locate the correct grammar from the enormous class of grammars that are largely (though not necessarily entirely) compatible with the primary linguistic data.
[12] Hornstein and Lightfoot (1981b). See also Hornstein, 1984, Chapter 1.
[13] See, for example, Lightfoot (1982) pp. 51–57, and the essays in Hornstein and Lightfoot (1981a).
[14] Actually, what is important here is not that all the correct grammars share certain properties, but only that they are all members of some quite restricted class. Since the distinction makes little difference to our current concerns, we shall ignore it in what follows.
[15] Lightfoot (1982), p. 21.
[16] Searle (1974), p. 22.
[17] Chomsky (1972), p. 88.
[18] For more on the contrast between discrete and distributed representations, see Ramsey, Stich and Garon (forthcoming).
[19] Prior to connectionism there were some dissenting voices. See, for example, Stich (1971), Cummins (1977) and Stabler (1983). However, a common response to the critics

was the question: 'What else could it be?' Thus, for example, Berwick writes, "I don't share Stabler's fear that 'we ought to worry about whether we can justify the current emphasis on program-using systems in theories about how people process language'. It's the only game in town" (1983, p. 403).

[20] Rumelhart, Hinton, and Williams (1986).

[21] For more on connectionist learning techniques, see Rumelhart and McClelland (1986a), Chapters 5, 7, 8, and 11. See also Hinton (1987).

[22] Chomsky (1980c), p. 238.

[23] Hanson and Kegl (1987).

[24] An auto-associator network is one that attempts to reproduce on the output nodes whatever input it receives on the input nodes. Hence, its input also serves as its teacher and source of the error signal during the training period.

[25] It should be noted here that Hanson and Kegl do not feel their model supports anti-nativist conclusions; rather, they believe it helps to delineate those aspects of grammatical structure which can be extracted from the data.

[26] Hanson and Kegl (1987), p. 106.

[27] Ibid., p. 117.

[28] For a critical analysis of this network, see Pinker and Prince (1988).

[29] Sejnowski and Rosenberg (1987); Gorman and Sejnowski (forthcoming).

[30] See Rumelhart, Hinton and McClelland (1986), p. 53

[31] See McClelland, Rumelhart and Hinton (1986), p. 43.

[32] Sampson (1987a), p. 877. Sampson (1987b), p. 643.

[33] See Goodman (1965). Compare Morgan (1986), p. 15: "It is fairly trivial to demonstrate that no unbiased inductive mechanism can reliably succeed in solving this sort of projection problem."

[34] Wexler and Culicover (1980), p. 10.

[35] We are indebted to Jeffrey Elman for convincing us of the importance of this point.

[36] Rumelhart and McClelland (1986c), p. 140.

[37] Rumelhart (personal communication).

REFERENCES

Berwick, R.: 1983, 'Using What You Know: A Computer–Science Perspective', *Behavioral and Brain Sciences* **6,** 402–403.

Charniak, E. and E. Santos: 1986, 'A Connectionist Context-free Parser which is not Context-free, but then it is not Really Connectionist Either', Department of Computer Science, Brown University.

Chomsky, N.: 1965, *Aspects of the Theory of Syntax*, MIT Press, Cambridge, Massachusetts.

Chomsky, N.: 1966, *Cartesian Linguistics: A Chapter in the History of Rationalistic Thought*, Harper and Row, New York.

Chomsky, N.: 1969, 'Comments on Harman's Reply', in S. Hook (ed.), *Language and Philosophy*, New York University Press, New York, pp. 152–159.

Chomsky, N.: 1972, *Language and Mind*, Harcourt Brace Jovanovich, New York.

Chomsky, N.: 1975, *Reflections on Language*, Pantheon Books, New York.

Chomsky, N.: 1980a, *Rules and Representations*, Columbia University Press, New York.
Chomsky, N.: 1980b, 'Rules and Representations', *Behavioral and Brain Sciences* **3**, 1–61.
Chomsky, N.: 1980c, 'Recent Contributions to the Theory of Innate Ideas: Summary of Oral Presentation', in H. Morick (ed.), *Challenges to Empiricism*, Hackett, Indianapolis, pp. 230–40.
Chomsky, N.: 1986, *Knowledge of Language*, Praeger, New York.
Chomsky, N.: 1988, *Language and Problems of Knowledge*, MIT Press, Cambridge, Massachusetts.
Cottrell, G.: 1985, 'Connectionist Parsing', in *Proceedings of the Seventh Annual Cognitive Science Society*, pp. 201–11.
Cummins, R.: 1977, 'Programs in the Explanation of Behavior', *Philosophy of Science* **44**, 269–87.
Elman, J.: 1988, 'Finding Structure in Time', CRL Technical Report 8801.
Fanty, M.: 1985, 'Context-Free Parsing in Connectionist Networks', Technical Report No. 174, Department of Computer Science, University of Rochester.
Fodor, J.: 1981, *Representations*, MIT Press, Cambridge, Massachusetts.
Gold, E. M.: 1967, 'Language Identification in the Limit', *Information and Control* **10**, 447–74.
Goodman, N.: 1965, *Fact, Fiction and Forecast*, 2nd ed., Bobbs Merrill, Indianapolis.
Gorman, R. and T. Sejnowski: forthcoming, 'Learned Classification of Sonar Targets Using a Massively Parallel Network', to appear in *IEEE Transactions: Acoustics, Speech, and Signal Processing*.
Hanson, S. and J. Kegl: 1987, 'PARSNIP: A Connectionist Network that Learns Natural Language Grammar from Exposure to Natural Language Sentences', in *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, pp. 106–19.
Hinton, G.: 1987, 'Connectionist Learning Procedures', Tech Report No. CMUCS-87-115.
Hornstein, N.: 1984, *Logic as Grammar*, MIT Press, Cambridge, Massachusetts.
Hornstein, N. and D. Lightfoot: 1981, *Explanations in Linguistics*, Longman, London.
Hornstein, N. and D. Lightfoot: 1981b, 'Introduction', in Hornstein and Lightfoot (1981a), pp. 9–31.
Lightfoot, D.: 1982, *The Language Lottery*, MIT/Bradford Press, Cambridge, Masssachusetts.
McClelland, J., D. Rumelhart, and D. Hinton: 1986, 'The Appeal of Parallel Distributed Processing', in Rumelhart and McClelland (1986a), Vol. I.
McClelland, J. L., and A. Kawamoto: 1986, 'Mechanisms of Sentence Processing: Assigning Roles to Constituents', in Rumelhart and McClelland (1986a), Vol. II.
Morgan, J.: 1986, *From Simple Input to Complex Grammar*, MIT Press, Cambridge, Massachusetts.
Papert, S.: 1988, 'One AI or Many?', *Daedalus* **117**, 1–14.
Peters, S.: 1972, 'The Projection Problem: How is a Grammar to be Selected?', in *Goals of Linguistic Theory*, Prentice-Hall, Englewood Cliffs, New Jersey.
Pinker, S., and A. Prince: forthcoming, 'On Language and Connectionism: Analysis of a Parallel Distributed Processing Model of Language Acquisition', to appear in *Cognition*.
Quine, W. V.: 1969, 'Linguistics and Philosophy', in S. Hook (ed.), *Language and Philosophy*, New York University Press, pp. 95–98.

Ramsey, W., D. Rumelhart, and S. Stich: forthcoming, *Philosophy and Connectionist Theory*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Ramsey, W., S. Stich and J. Garon: forthcoming, 'Connectionism, Eliminativism and the Future of Folk Psychology', to appear in Ramsey, Rumelhart, and Stich.

Rumelhart, D., G. Hinton and J. McClelland: 1986, A General Framework for Parallel Distributed Processing', in Rumelhart and McClelland (1986a), Vol. I.

Rumelhart, D., G. Hinton and R. Williams: 1986, 'Learning Internal Representations by Error Propagation', in Rumelhart and McClelland (1986a), Vol. I.

Rumelhart, D. and J. McClelland: 1986a, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vols. I & II, MIT/Bradford Press, Cambridge, Massachusetts.

Rumelhart, D. and J. McClelland: 1986b, 'On Learning the Past Tense of English Verbs', in Rumelhart and McClelland (1986a), Vol. II.

Rumelhart, D. and J. McClelland: 1986c, 'PDP Models and General Issues in Cognitive Science', in Rumelhart and McClelland (1986a), Vol. I.

Sampson, G.: 1987a, 'Review Article. Parallel Distributed Processing', *Language* **63,** 871–86.

Sampson, G.: 1987b, 'A Turning Point in Linguistics', *Times Literary Supplement*, June 12, p. 643.

Searle, J.: 1974, 'Chomsky's Revolution in Linguistics', in Gilbert Harman (ed.), *On Noam Chomsky: Critical Essays*, Doubleday, New York, pp. 2–33.

Sejnowski, T. C. Rosenberg 1987, 'Parallel Networks that Learn to Pronounce English Text', *Complex Systems* **1,** 145–68.

Selman, B. and G. Hirst: 1985, 'A Rule-Based Connectionist Parsing System', *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*.

St. John, M. F. and J. L. McClelland: 1988, 'Learning and Applying Contextual Constraints in Sentence Comprehension', in *Proceedings of the 10th Annual Cognitive Science Society Conference*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Stabler, E.: 1983, 'How Are Grammars Represented?', *Behavioral and Brain Sciences* **6,** 391–421.

Stich, S.: 1971, 'What Every Speaker Knows', *Philosophical Review* **80**, 476–96.

Stich, S.: forthcoming, 'The Dispute Over Innate Ideas', to appear in M. Dascal et al. (eds.), *Sprachphilosophie: Ein Internationales Handbuch Zeitgenossischer Forschung*.

Waltz, D. L. and J. B. Pollack: 1985, 'Massively Parallel Parsing: A Strongly Interactive Model of Natural Interpretation', *Cognitive Science* **9,** 51-74.

Wexler, K. and P. W. Culicover: 1980, *Formal Principles of Language Acquisition*, MIT Press, Cambridge, MA.

Dept. of Philosophy
University of Notre Dame
Notre Dame, 46556
U.S.A.

and

Dept. of Philosophy
Rutgers, The State University of New Jersey
New Brunswick, NJ 08903
U.S.A.