4

Connectionism without Tears

Mark S. Seidenberg

What accounts for the cool reaction to the emergence of connectionism in the 1980s on the part of people who study language for a living? Most of the critical assessments of connectionism that followed the initial explosion of interest in the approach came from people such as Bever, Fodor, Pinker, and Prince, whose works are firmly situated within the mainstream of linguistics and psycholinguistics. It will be an interesting project for a future student of the history of ideas or the sociology of science to investigate why, as Prince and Smolensky recently observed, connectionism was seen as 'at best orthogonal and at worst antithetical to the goals of linguistic theory.'[1] This issue is of particular interest to me as someone who was trained in the standard linguistic-psycholinguistic school of thought but has utilized connectionist modelling techniques in studies of language processing. My own work has been grounded in the belief that connectionism and linguistics have more in common than some of the more polemical critiques of the approach would suggest. In this paper I will discuss some of the issues that have tended to separate the two approaches and describe some potentially interesting points of contact.

Nativism and Empiricism

It is clear that many linguists view connectionism as a revival of the radical empiricist approach that dominated the dark ages in psychology–the behaviourist era. Pinker, in particular, equates connectionism with 'associationism,' a move that has the effect of eliciting, almost by reflex, the intellectual and emotional antipathy that most linguists feel towards the behaviourist account of language (establishing a kind of guilt by association(ism)).[2] Rumelhart and McClelland's (1986) claims as to what their model of past tense learning showed about the acquisition of language must surely have elicited widespread feelings of deja vu all over again. I think that the attempt to equate connectionism with radical empiricism is a mistake. The correct point of contact is not with the learning theories developed by the animal behaviourists of the 1950s but, rather, with the learnability approach developed by the linguists and psycholinguists of the 1970s and 1980s (for example, Wexler and Culicover 1980; Pinker 1979; Baker and McCarthy 1981).[3] Learnability theory attempts to explain language acquisition in terms of several factors: the initial state of the organism (i.e., innate capacities that are probably species- and domain-specific), the steady state to be achieved (i.e., grammar), the input to the child ('Motherese' and other aspects of the environment; Newport et al. 1977), and the child's capacity to learn. The same factors govern the operation of connectionist models. The initial state of the organism can be equated with the initial configuration or architecture of the model. Steady-state behaviour represents the target to which the model should converge. The model's behaviour changes in response to experience – the 'input to the child.' What connectionism adds to the learnability approach is a novel way of representing knowledge and a substantive theory of learning. The novel way of representing knowledge is in terms of weights on connections between units. The substantive theory of learning is given by the many learning algorithms that operate over networks employing these distributed representations. The main implication I draw from explorations of learning algorithms such as back-propagation (Rumelhart, Hinton, and Williams 1986) is that far richer underlying structures can be recovered from far noisier data than anyone ever imagined. (Indeed, the reinforcement learning techniques of the behaviourists occupy a small and undistinguished corner in a very large space of learning algorithms.) Applications of such learning principles to human behaviour (for example, Gluck and Bower 1988; Elman 1990; Hare 1990) suggest that they capture important aspects of at least some ways in which people learn.

At the same time, it is obvious that the behaviour of connectionist systems is highly dependent on their initial configuration. To the extent that there is evidence that innate capacities govern the acquisition of language (and I think there is), they can be straightforwardly incorporated in connectionist models. McClelland and I provided a simple illustration of this point with our model of word recognition and pronunciation (Seidenberg and McClelland 1989). We described a simple multi-layer backprop net that was trained on a large corpus of monosyllabic words and which simulated numerous aspects of peoples' performance in behavioural experiments employing such stimuli (subjects in these experiments perform tasks such as reading words aloud or deciding whether or not strings of letters form words). The model that closely simulated many aspects of skilled performance was configured with 200 hidden units. We also replicated this simulation using an architecture that provided only 100 hidden units. The disabled model could master some aspects of word recognition and naming but made systematic errors. The interesting part of this exercise was that the pattern of performance exhibited by the disabled model resembled that of some children who are dyslexic (i.e., fail to acquire age-appropriate reading skills). Thus, whether the model performed like a skilled reader or a dyslexic reader depended on its initial configuration – specifically, whether it contained sufficient units and connections to encode all of the information that the task demanded. One way to view these results is that humans have an innate capacity to allocate different neural resources to different tasks, such as reading or speaking (subject to considerable plasticity; Neville 1988). For unknown reasons, dyslexic children may dedicate too little in the way of neural resources to the task of learning to read. Though speculative, this theory is suggestive in light of recent evidence concerning morphological anomalies in the brains of dyslexic children revealed by magnetic resonance imaging (Hynd and Semrud-Clikeman 1989). In sum, the model's performance is both consistent with and lends substance to the idea that tasks such as learning to read are subject to biological constraints.

I realize that when linguists talk about innate capacities for language they have in mind something more specific than a tendency for certain brain areas to be recruited for certain tasks. The inventory of hypothesized innate capacities includes language-specific knowledge structures, tendencies to analyse linguistic input in specific ways, and constraints on the range of hypotheses that are formed, among others. Connectionism is equally compatible with these ideas. Moreover, it provides a basis for exploring exactly how innate capacities of various sorts would affect the course of acquisition. For example, it should be possible to determine why only certain types of generalizations are formed given the starting configuration of the system, the input to the child, and a specific learning algorithm. Again, our model provides a simple illustration. The model addresses a certain range of phenomena concerning word recognition. Our principal focus was on the acquisition of knowledge concerning the correspondences between spelling and sound. The goal was to understand the kinds of generalizations concerning these correspondences that could be learned on the basis of experience. In its initial configuration the model was endowed with ways of representing orthography and phonology. Although the correspondences between the codes were learned, the codes themselves were not. Thus the model tacitly embodies the idea that children who are learning to read already possess considerable knowledge of the sound structure of the language (for example, its phonemic inventory and phonotactic constraints). Some of this knowledge is thought to derive from innate capacities to analyse speech in special ways (for example, Liberman and Mattingly 1985). In an admittedly oversimplified way (the phonological representation we used was, after all, Wickelphones), the model captures this idea. Our particular model did not address where these phonological representations come from because that was not its primary focus. Assume, for the sake of argument, that these aspects of phonology are entirely innate. The model could then be taken as having exemplified how biological constraints of a certain sort constrain what can be learned. The general point is that if, in fact, it is the case that knowledge of some kinds of phonological distinctions (or of other aspects of language) is innate, this can be represented in a net. It would certainly be a reasonable goal, for the future, to endow such models with exactly the innate capacities for which there is independent evidence.

These observations merely establish the simple point that connectionism is compatible with nativism. More important, however, I think that rather than being merely compatible with the nativist view, connectionism is likely to provide what is needed in order to establish the essential correctness of this view with regard to language. The learnability equation I gave above seems to be the proper way to decompose the language acquisition problem but, unfortunately, it yields an equation with more than one unknown. Rather more attention has been focused on characterizing the steady state than on understanding the learning component or the input to the child. Nonetheless, very strong inferences about the initial state of the organism have been drawn. This neglect of the role of learning is perhaps understandable given the meager kinds of learning principles available during much of the history of generative grammar. Lack of interest in learning may also have contributed to a lack of attention to many aspects of the child's experience. Connectionism now provides an interesting theory of learning, motivating empirical studies of whether children utilize such principles, and closer assessments of the behavioural input relevant to the acquisition process. The important implication is that, with a serious theory of human learning nearly in hand, a better understanding of the nature of the input to the child (for example, Fernald 1984; Hirsch-Pasek, Treiman, and Schneiderman 1984), and a rigorous theory of the structure of language (for example, Chomsky 1981), it may be possible at last to solve for the final unknown in the equation – the biological endowment of the child relevant to language. That there is such an endowment seems obvious to me (see Seidenberg 1985a; Seidenberg and Petitto 1987) but it is hard to be specific about exactly what is innate without at the same time knowing what can be learned.

It is an empirical question, of course, but studies of connectionist learning suggest the possibility (in my view, the likelihood) that more of the task of acquiring a language is accomplished by means of learning principles operating on relatively fragmentary, noisy input than was previously assumed. Although the same learnability equation is involved, the division of labour among the various components may turn out to be somewhat different than standard accounts suggest. It would be a serious error to underestimate the power and importance of these learning principles. They will force the reassessment of 'poverty of the stimulus' arguments; puzzles that currently occupy child language researchers, such as how the child escapes from incorrect generalizations working only from positive exemplars, will disappear (connectionist models already learn without direct negative evidence). One of the regrettable consequences of the contentious way in which connectionism was presented to linguists (for example, Rumelhart and McClelland 1986) and the strenuous way in which it was attacked (for example, Pinker and Prince 1988) is that these points of contact between connectionism and learnability theory were obscured. It would be incorrect to conclude that the views expressed by Rumelhart and McClelland in their 1986 paper exhaust the range of possibilities afforded by the broader connectionist framework, though I think that is what in fact happened. Fortunately – in my view, it was inevitable – we are beginning to see the appearance of research that exhibits an appreciation of both the value of connectionism and of the facts about child language that need to be explained (for example, Punkett and Marchman 1991). I believe that as the polemics recede into the background, insights from connectionism are likely to be absorbed by more of the people who study language acquisition, with profound impact on our understanding of the phenomenon.

Connectionism and Linguistic Explanation

As Prince and Smolensky suggest, the initial work that emerged from the connectionist framework could be seen as largely orthogonal to the interests of theoretical linguists. One influential perspective was articulated by Pinker and Prince (1988). They take the view that knowledge of a language consists of knowledge of different kinds of systems of rules. According to this view, the task of the linguist is to identify the rules; the task of the child, to acquire them. Pinker and Prince observed that connectionist models such as the one proposed by Rumelhart and McClelland (1986) fail to capture the rule-governed character of human languages. The inadequacies of the Rumelhart and McClelland proposal about the English past tense led them to question whether connectionist models could contribute to understanding other linguistic phenomena. At best, they suggested, connectionist models might implement the kinds of rule systems posited within linguistic theory. Implementing rule systems in connectionist hardware might be a useful thing to do – it might contribute to understanding how the rules are realized in the brain, for example – but in their view the important generalizations are captured at the level of the rules themselves.

In light of the recent history of syntactic theory, it is ironic that the debate about connectionism and symbolic systems in general, and about connectionism and language in particular, was framed in terms of the issue of rules. Whereas earlier theories (for example, Chomsky 1965) incorporated numerous structure-specific rules, current theories (for example, Chomsky 1981) do not. Government-binding theory is principle-based, not rule-based; a sentence is well-formed if it satisfies the constraints imposed by the several modules in the grammar. In a sense, well-formedness is treated as a constraint satisfaction problem (Stabler 1991), which is certainly congenial to the connectionist approach. Similarly, in current theories, language acquisition is not seen as the process of acquiring rules but, rather, of setting parameters governing the range of possibilities afforded by universal grammar. The irony, then, is that Pinker and Prince based their critique of connectionist accounts of language on a notion of rule that has little relevance to the dominant theory in syntax, which has provided the intellectual core of theoretical linguistics.

I am not prepared to speculate about the potential for convergence between connectionism and syntactic theory; for some interesting, though preliminary, work that is relevant to this issue see Berg (1991), who describes a recurrent net that learns aspects of X-bar syntax. Leaving syntax aside, the view of linguistic theory offered by Pinker and Prince is still widely held in areas such as phonology and morphology, and it is to these areas that I now turn. The fact that current syntactic theory largely eschews the notion of rule (and especially rule-ordering) compelled Bromberger and Halle (1989) to defend the proposition that phonology is really different (insofar as it demands the use of these formal mechanisms). Every phonological theory of which I am aware (autosegmental phonology; metrical phonology; lexical phonology) follows the Bromberger-Halle (Pinker-Prince) line about rules, though they differ, of course, in terms of the types of rules allowed as well as in other respects. Morphological theories typically embrace this approach as well (but see Bybee 1985, for a somewhat different view).

The picture that Pinker and Prince draw is a tidy one. Linguists have developed theories (for example, of phonology and morphology) employing certain kinds of rules; these theories are demonstrably correct insofar as they capture important generalizations that would otherwise be unexplained. This leaves connectionism with a dilemma: either connectionist theories are wrong, insofar as they behave in ways that are inconsistent with the notion of rule relevant to linguistic theory, or they exhibit rule-governed behaviour, in which case they are mere implementations.

I think that connectionism has more to contribute to understanding language than the circumscribed role that Pinker and Prince assign to it, and in the material that follows I attempt to establish why. Since I cannot address all aspects of phonology and morphology for both practical and intellectual reasons, I will follow Pinker and Prince's lead and focus on a narrow but interesting set of phenomena: the past tense of verbs in English – the same phenomena that were at the center of their celebrated critique. It might first be observed, however, that morphological theory in its current state is nowhere near as tidy as Pinker and Prince imply. Morphology is probably the least developed of all the major subareas of linguistic theory. Aside from the fact that there is no unifying theoretical framework, there is debate about the range of phenomena that constitute the subject matter of the field. Some aspects of traditional morphology have been absorbed by phonology (for example, Kiparsky 1982), others by syntax (Selkirk 1984), leaving undecided whether morphology has a subject matter of its own and what its boundaries are (Anderson 1982). The disagreements here run very deep, and there are several competing theoretical frameworks.

English inflectional morphology (which includes past tense formation) is a rather simple system, and it might therefore be thought that, whatever the theoretical debates about, for example, triliteral roots in Arabic, there would be broad agreement about it. Such is not the case. There is a general commitment to generating at least some past tense forms by rule but exactly how many and what types of rules are involved varies across theories. A common (though by no means universal) assumption is that the lexical component of the grammar contains only a listing of idiosyncratic forms, such as irregular past tenses (so-called strong alternations such as BRING-BROUGHT or SING-SANG). However, which past tenses are irregular, and whether the irregularities are morphological or phonological, are unclear. The problem is illustrated by the fact that whereas Pinker and Prince consider alternations such as THINK-THOUGHT and SING-SANG to be idiosyncracies to be listed in the lexicon, Halle and Mohanan (1985) derive them by rule (THINK-THOUGHT, for example, is derived by a phonological rule the only other application of which is to BRING to form BROUGHT). In fact, Halle and Mohanan (1985) derive almost all strong verbs in English by rule. And theirs is not the only product in the marketplace. There have been several theoretical analyses of inflection, including past tense morphology, in the recent past and no one of them has come to predominate. Even the basic nature of the phenomenon is undecided: there are fundamental disputes as to whether the past tense and other aspects of inflection should be treated as part of an autonomous morphological component, as phonological phenomena, or as part of syntax (Anderson 1982; Spencer 1991).

In short, current treatments of inflectional morphology admit many possibilities other than the view promoted by Pinker and Prince. The lack of consensus about foundational issues, such as scope of morphological theory, calls into question their contention that existing linguistic theories provide adequate accounts of the phenomena, which connectionists could only aspire to implement.

Can a single mechanism accommodate both rules and exceptions?

Ignoring the technical disputes that animate current discussions of inflectional morphology, Pinker and Prince (1988) present an appealingly simple view: regular past tense forms are generated by rule; irregular forms are listed in the lexicon. The fact that identifying which forms are irregular is itself a theoretical problem, and the fact that at least some theories hold that even regular, rule-governed forms (for example, familiar ones) are listed (Halle 1973) does not figure in their discussion. They present a generic framework much like the one developed by Aronoff (1976). Aronoff's model takes the word as the primary unit of morphological analysis, in contrast to theories based on other units, for example, morphemes (Halle 1973). Words that are formed by entirely regular, productive processes are not listed in the lexicon. Pinker and Prince discuss a range of facts which they take to support the general distinction between past tense forms that are generated by rule and those that must be listed as exceptions. For example, the irregular pasts include neighbourhoods of phonologically-similar pairs such as SING-SANG and RING-RANG; the regular forms do not exhibit this clustering, because the rules can apply to any present tense form without regard for its phonological composition. The McClelland and Rumelhart model does not enforce this distinction and therefore, it is argued, cannot capture these systematic differences between regular and irregular forms. This raises questions as to whether any connectionist system could do better.

The general approach that Pinker and Prince present is not limited to the past tense; there have been many attempts to characterize various aspects of linguistic knowledge in terms of rules. The problem for such theories is what to do with cases in which the rules fail. These cases seem endemic to human language. Consider some examples. The written form of English is alphabetic; hence there are systematic correspondences between spelling and pronunciation. It is often assumed that these correspondences can be formulated as rules; see Hanna, Hanna, Hodges, and Rudortf (1966), Wijk (1966) and Venezky (1970) for attempts to list them. That the system is rule-governed seems to be indicated by patterns such as MINT-LINT-SPLINT-HINT; that people know such rules seems indicated by the fact that they can pronounce novel strings such as BINT. On this view, then, one of the child's first tasks in learning to read is to master these rules. Although (as in morphology) there is no generally agreed-upon list of rules, it is clear that there will be numerous exceptions to them, however they are formulated. What kind of rules would accommodate minimal pairs such as MINT/PINT, GAVE/HAVE, PAID/SAID, LEAF/DEAF and triples such as POSE/DOSE/LOSE or COUGH/DOUGH/PLOUGH? Presumably items such as COLONEL, CORPS, and ONCE will be treated as exceptions to any felicitous set of rules. Thus, the spelling-sound correspondences of English are apparently rule-governed, but the system admits many exceptions. As in the case of past tense inflectional morphology, there are disagreements about the exact content of the rules, yielding uncertainty as to which items are exceptions. For example, should DONE be listed as an exception or generated by a minor rule that also applies to NONE (analogous to Halle and Mohanan's rule that only applies to BRING-BROUGHT and THINK-THOUGHT)?

The mapping between spelling and syllabification is another domain that has received a rules plus exceptions analysis. Hansen and Rodgers (1968) developed a set of rules for syllabifying written English words on a strictly orthographic basis; these rules were later incorporated by Spoehr and Smith (1973) into a theory of visual word recognition. The rules work correctly in many cases, and the Spoehr and Smith research suggests that they capture something about the way people process words in reading. There again are cases where the rules fail, however, illustrated by minimal pairs such as BAKED-NAKED, DIES-DIET, and WAIVE-NAIVE.

Similar phenomena occur in other areas. As I have noted, the standard approach is to treat inflectional morphology as rule-governed but, as Spencer (1991) notes, 'Inflectional morphology is notorious for being morphologically idiosyncratic.' In English, of course, there is the past tense, typically formed by adding /d/. This system is overwhelmingly regular: there are about 4400 verbs in the Francis and Kucera (1982) count, of which perhaps five percent are irregular forms such as RUN-RAN or TAKE-TOOK. The irregular cases tend to cluster among the higher frequency words in the language (see Seidenberg 1989, for a discussion of why they do); hence, on a token-wise basis, the percentage of irregular forms is somewhat greater, though still less than for the regular forms. In short, there are both rule-governed cases and exceptions, with the former greatly outnumbering the latter. The situation is similar with respect to forming the plural. There is a rule, add -s, as in PAN-PANS or FAN-FANS; there are exceptions such MAN-MEN and SHEEP-SHEEP (I am suppressing irrelevant details here concerning, for example, the conditions that determine whether the -s is realized as /s /,/z/, or /iz/).

Whereas inflectional morphology is relatively impoverished in English, derivational morphology is not. Many morphologists take as their goal the formulation of rules that will account for facts such as the following. A HEADACHE is a kind of ache in your head; a HEAD-COLD is a kind of cold in your head; a HEADLINE, however, is not a line in your head (that is a wrinkle). A DEADHEAD (in one sense) is a person who likes the Grateful Dead; a POTHEAD is a person who likes marijuana (and often Potheads and Deadheads are the same people). An EGGHEAD, however, is not a person who likes eggs, a BEACH-HEAD is not a surfer, and a BLACKHEAD is a kind of facial blemish. Although it is possible to formulate rules governing the formation of compounds in English (see Selkirk 1982 for discussion), it is doubtful whether they could be formulated in such a way as to cover all cases without admitting any exceptions.

Pinker and Prince's (1988) response to these sorts of phenomena is to suggest that the goal of a proper theory (for example, of verb morphology) should be the formulation of a set of rules that captures significant generalizations. Properties of language such as the ones sketched in the examples above seem to ensure that the kind of rules to which they are committed will necessarily fail in many cases (the only way to avoid this would be to have rules that apply to individual items, which would trivialize the notion of rule). The proposal for what to do with these anomalies is simply to list them separately. These are the items that will have to be learned 'by rote.' Thus, the idea that language is rule-governed at various levels of structure (exemplified in the Pinker and Prince paper by the treatment of past tense morphology) is preserved by introducing a second descriptive mechanism – a list – to deal with cases in which the rules fail.

It seems to me that any system can be described by a set of rules if the rules do not have to apply to all cases. Pinker and Prince claim as a major discovery of linguistic theory that languages are rule-governed at different levels; however, it is hard to see how any other outcome would have been possible, given their notion of rule and the existence of a second means of accommodating all of the cases where the rules fail. This is like saying that all of the observations in my experiment fit a particular hypothesis except for the ones that I have decided to exclude. Aside from the fact that it seems a logical necessity that any set of phenomena can be partitioned in this way, Pinker and Prince assert that this rules-and-exceptions approach accounts for facts about verb morphology that would otherwise be unexplained. For example, it is thought to explain why clusters of phonologically-similar past tenses (RING-RANG, SING-SANG, etc.) only occur among the exceptions. Again, however, it is hard to see how any other outcome could obtain. That such similarity clusters as RING-RANG/SING-SANG exist is an interesting fact (one that is itself in need of explanation, since it is easy to imagine a system in which this patterning does not occur) but, given their definition of the rule, it seems tautological that if such patterning does occur it must be confined to the exceptions.

These logical considerations aside, it could be the case that the assertion that knowledge consists of a set of rules and a list of exceptions is merely true. Here it is worth returning to the area in which this idea has been investigated most thoroughly, the mapping between spelling and pronunciation. In so-called dual-route models of reading aloud, there are two types of knowledge representation: a set of rules governing spelling-sound correspondences (sometimes termed grapheme-phoneme correspondence rules) and a lexicon in which the irregular cases must be listed (Coltheart 1978). There are also two pronunciation mechanisms: the pronunciations of rule-governed items are generated by applying the rules; the pronunciations of the words that violate the rules are looked up in the list.

Dual-route models have been justified on a number of grounds (see Carr and Pollatsek 1985 and Seidenberg 1985b for reviews). They developed in response to a variety of empirical phenomena concerning reading aloud and the acquisition of this skill; they have also provided a useful framework for studying some kinds of reading impairment that occur as a consequence of brain injury (Patterson, Marshall, and Coltheart 1985). These models have also been justified on the basis of the intuition that no single type of knowledge representation or process could simultaneously account for the ability to read rule-governed items such as GAVE, irregular items such as HAVE, and novel, nonword items such as MAVE. Elsewhere I have termed this the 'central dogma' of dual-route models (Seidenberg 1988). This view was put forward with admirable clarity by Coltheart (1987):

> A crucial implication of this distinction [between the two pronunciation mechanisms], an implication around which much work on normal and abnormal reading has been organized, is that the two procedures are not capable of producing correct responses for every type of orthographic input ... The word-level procedure allows correct reading aloud only when the orthographic stimulus is a word .... In contrast the subword level procedure guarantees correct reading aloud only when the orthographic stimulus is a

> regularly spelled word or a nonword. According to this general approach to modelling oral reading, then, correct reading of nonwords requires a procedure for subword-level translation from orthography to phonology, whereas correct reading of words irregular in spelling-sound correspondences requires a procedure for word-level translation (xvi).

There are now two connectionist models that directly contradict this central dogma. Both the Sejnowski and Rosenberg (1986) NETtalk model and the Seidenberg and McClelland model generate phonological codes for words on the basis of orthographic input. The models generate correct output for both 'rule-governed' cases such as LIKE and TAKE and irregular cases such as HAVE and GIVE. The important point is that within the limited domains in which these models operate (for example, in the Seidenberg and McClelland model, the domain is monosyllabic words), they illustrate the claim that connectionist nets can generate both rule-governed cases and exceptions by means of a single mechanism. These models appear to refute the central dogma as it applies to spelling-sound knowledge.

The next question is whether a similar model could successfully accommodate the past tense. One of the central claims of McClelland and Rumelhart (1986) is that both rule-governed instances and exceptions can be derived from a single underlying system of units and connections in learning the past tense. Pinker and Prince are correct in suggesting that the McClelland and Rumelhart model of the acquisition of the past tense does not substantiate this claim. As an account of an aspect of child language (as opposed to a demonstration of some interesting properties of connectionist networks), the model is fatally flawed. Leaving aside this particularly unhappy case, it could be asked whether other attempts might be more successful in refuting the central dogma as it applies to past tense acquisition.

The answer to this question is as yet unknown and will not be known until someone develops a model that addresses the many important empirical phenomena described in the Pinker and Prince paper and other phenomena as well (see below). The many similarities between the system of spelling-sound correspondences and the past tense in English are certainly suggestive, however. Both systems are 'rule-governed' but admit many exceptions; in both cases the exceptions tend to cluster among the higher frequency words in the language and thus are overrepresented among the words to which the child is first exposed. In fact, the spelling-sound system exhibits all of the differences between rule-governed cases and exceptions that Pinker and Prince (1989:188) ascribe to verbs:

(a) Irregular verbs cluster into "family resemblance groups" that are phonologically similar: BLOW/BLEW; GROW/GREW; THROW/THREW.' Irregularly pronounced words also cluster: DONE/NONE; PUSH/BUSH; BREAK/STEAK.

(b) Irregular pasts can be fuzzy in their naturalness of acceptability ... In contrast, regular verbs, unless they are similar to an irregular cluster, have no gradient of acceptability based on their phonology.' The spelling-sound correspondences of COLONEL, ACHE, and BREAST seem less natural than those of KERNEL, TAKE, and BEAST.

(c) There are no sufficient conditions for a verb to be in any irregular class: though BLOW becomes BLEW in the past, FLOW becomes FLOWED; though RING becomes RANG, STRING becomes STRUNG and BRING becomes BROUGHT. In contrast, a sufficient condition for a verb to be regular is that it not be irregular.' HAVE-GAVE; SAID-PAID; BONE-DONE-GONE.

(d) Most of the irregular alternations can only apply to verbs with a certain structure: the pattern in 'send/sent,' namely to change a *d* to a *t*, requires that there be a *d* in the stem to being with. The regular rule, which adds a *-d* to the stem, regardless of what the stem is, can cover all possible cases by its very nature.' I am not clear what is being claimed here, other than that the exceptions are idiosyncratic patterns. However, the pattern -AVE requires the letter H in the initial position in order to change to /av/and similarly for other cases.

In summary, I hope to have established that, inadequacies of the Rumelhart and McClelland verb learning model aside, the idea that a single system might be responsible for both rule-governed items and exceptions is quite viable; in at least one domain there are existing models that implement the idea. It remains to be determined whether a similar account will apply to the past tense. However, the kinds of phenomena that Pinker and Prince take to implicate very different types of knowledge representations and processing mechanisms for rule-governed items and exceptions, and to be incompatible with connectionist models, also occur in the domain of spelling-sound correspondences, for which plausible models already exist. A number of people who attended closely to Pinker and Prince's critique of the Rumelhart and McClelland model have begun developing successors to it (for example, Plunkett and Marchman 1991; MacWhinney in press;

Cottrell and Plunkett 1991). It therefore seems likely that we will soon have a clearer picture of the relevance of connectionist models to verb learning. Although none of the existing models as yet achieves descriptive adequacy, a number of interesting results have been achieved. Both the models of MacWhinney (in press) and Cottrell and Plunkett (1991) use a single network to generate both regular and exceptional past tenses. It will be interesting to determine whether extensions of these models or others like them will be able to accommodate the entire range of facts.

Are rules and exceptions sufficient?

To this point I have argued that systems of knowledge that have the character of past tense morphology in English are compatible with known properties of simple connectionist networks. In light of the controversy that followed the Rumelhart and McClelland (1986) model I should stress that in the absence of an adequate, implemented model these observations are merely suggestive. For the sake of argument, let us assume that an adequate a model could be constructed. Two related questions then arise. One concerns whether such a net would merely implement the two mechanisms that Pinker and Prince envision. For example, the network might partition itself so that some units and connections are dedicated to implementing the rules and others to implementing the list of exceptions. The second question concerns whether there would be any way to determine which is the correct account. In this section I will present evidence suggesting that peoples' behaviour departs from that which would be expected if their knowledge of the past tense were represented in terms of rules. Moreover, I will argue that these departures from orderly rule-governed behaviour can be understood in terms of simple properties of connectionist networks. Therefore, the connectionist approach is to be preferred because it captures generalizations that the rules-and-exceptions approache misses. These considerations also suggest a somewhat different relationship between connectionism and linguistic theory than Pinker and Prince's 'implementational' view.

As noted above, Pinker and Prince's view of the past tense is that it involves two types of knowledge: a set of rules and a list of exceptions. These, in turn, involve two types of learning: inducing a set of rules and learning the exceptions 'by rote.' Prasada, Pinker, and Snyder (1990) further assume that these types of knowledge entail different processing mechanisms: the rule-governed cases are generated; the exceptions are looked-up from storage in memory. They present the results of reaction time experiments in which subjects generated past tense forms aloud. The data were seen as supporting the distinction between generating output by rule versus lexical lookup, and they offer logical arguments suggesting that the phenomena cannot be properly understood without making this distinction. This is the central dogma that I questioned previously, that is, that no single type of knowledge representation or process can simultaneously handle both the rule-governed cases and exceptions.

Models such as Seidenberg and McClelland's (1989) inspire the following alternative account. Knowledge of the past tense is encoded by weights on connections between units representing different types of knowledge (orthographic, phonological, grammatical, etc.). Learning involves adjusting the weights on the basis of experience. All forms are generated by a single process (for example, taking the present tense stem as input, along with other relevant information such as meaning, and computing the past tense as output). There are no rules in this type of model and no listing of irregular forms. The idea of a list of lexical entries is especially incongruent with this approach; in models of the lexicon employing distributed representations, there are no units or pools of units dedicated to individual words (see also Hinton 1986). Each word form is represented by a pattern of activation over one or more sets of units; each unit participates in the representation of many words.

As I have noted, the same theoretical alternatives arose in connection with spelling-sound knowledge. In the latter domain, however, there have been many studies of subjects' use of this knowledge under various conditions. There is a large body of data that is quite revealing about how this knowledge is represented and used, and it is sufficient to strongly call into question the dual-route account. In reading, the adequacy of the dual-route model began to be questioned because of the discovery of some unexpected phenomena that have come to be called consistency effects. The seminal study was by Glushko (1979). In the dual-route model, a word such as MUST is rule-governed and HAVE is an exception. Glushko asked a deceptively simple question: what about words such as GAVE? Under any plausible construal of the notion, GAVE is rule-governed. The rule presumably applies productively to SAVE, PAVE, RAVE and other words and can be used to pronounce nonsense words such as MAVE. However, the pronunciation of the -AVE pattern is inconsistent, owing to the irregular neighbour HAVE. In Glushko's experiments, subjects read such words aloud and their responses were timed. Subjects who are skilled, college-student readers perform this task at a high level of accuracy. Glushko replicated the earlier finding that irregular words (such as HAVE or SAID) take longer to pronounce than do regular words (such as MUST or LAKE). Unexpectedly, however, he found that so-called inconsistent words such as

GAVE or PAID also took significantly longer to read aloud than did entirely regular words. These results are important because both the regular and inconsistent words are rule-governed according to the standard dual-route approach. If such items were pronounced by applying rules, the two classes should have behaved alike. However, the inconsistent words yielded longer latencies due to the irregular neighbours.

Subsequent studies replicated the basic consistency effect but clarified it in a number of respects (see Jared, McRae, and Seidenberg 1991 for review). First, the effect is correlated with frequency: lower frequency words show larger effects, and among the highest frequency words in the language the effects are small or nonexistent. Second, the effects depend on reading skill: faster, more skilled readers show smaller effects. Third, the magnitude of the effect depends on the ratio of a word's friends and enemies. The friends of GAVE, for example, are all the rhyming -AVE words. Its only enemy is HAVE. Thus, regular words have many friends but no enemies; exception words have few friends and many enemies; and inconsistent words fall in-between. Jared et al. (1990) showed that the ratio of friends and enemies accounts for the results of about fifteen studies of consistency effects in the literature.

The important conclusion to be drawn from this research is that the generalization that accounts for word naming latencies is not whether they are rule-governed or exceptional. Rather, the correct generalization concerns the degree of consistency exhibited across a neighbourhood of similarly spelled items. The standard dual-route account suggests that the latency to pronounce a word should only depend on properties of the word itself: its frequency, length, and whether it is rule-governed or irregular. The empirical studies show that this assumption is false. Latencies systematically depend not only on the properties of the word itself but also on their neighbours. Thus, mechanisms for generating the pronunciations of words must take into account these relations among words.

The Seidenberg and McClelland (1989) model simulates these effects quite closely. The model was trained on a corpus of 2897 monosyllabic words, including almost all of the words used in studies of consistency effects. Hence it is possible to simulate each experiment using the same items as in the study. The naming latencies of the subjects are compared to an error score that is a quantitative measure of the model's performance. The fit between mean naming latencies and error scores for the same items is typically very good. Figure 4.1 provides a summary of the results of a study by Jared et al. (1990). The stimuli in the experiment were lower frequency inconsistent words (such as TINT, which is inconsistent because of PINT) and two control groups of matched regular words. All of the stimuli would be considered rule-governed in the

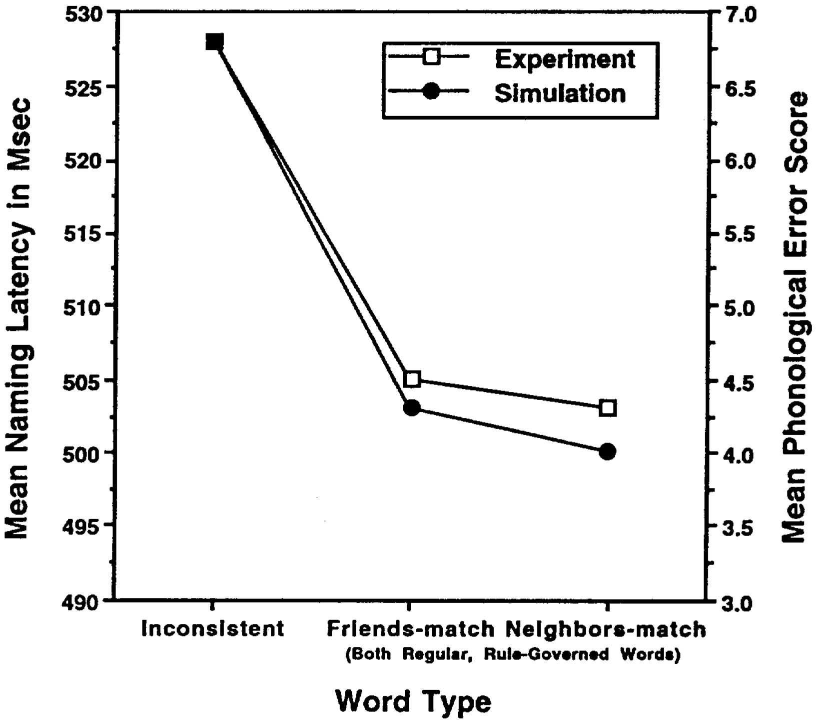

Figure 4.1: Results of an experiment on consistency effects (Jared et al. 1990). The 'friends match' and 'neighbours match' are two groups of entirely regular, rule-governed items.

dual-route account; however, the inconsistent items, which have enemies, yielded significantly longer naming latencies than did the regular words. As the figure indicates, the model produced very similar results. Many other simulations of this type are reported in Seidenberg and McClelland (1989) and Jared et al. (1990).

The explanation for why the model performs this way is simple. The weights mediating the computation from orthography to phonology encode facts about the frequency and consistency of spelling-sound correspondences in the lexicon. The model accounts for effects of lexical frequency (for example, McRae, Jared, and Seidenberg 1990) because frequency determines how often a word is presented during the training phase; words that are presented more often have a bigger impact on the weights. For the same reason, the model performs better on words containing sublexical spelling patterns that occur in many words. Thus, the model performs better on words containing spelling patterns that are consistently associated with a single pronunciation (for example, -UST in MUST, -IKE in LIKE) compared to inconsistent patterns associated with more than one pronunciation (for example, -OWN in TOWN, BLOWN). These outcomes are simply a consequence of how the learning algorithm operates given a significant fragment of the English lexicon with which to work. Multiple exposures to consistent patterns such as -UST push the weights towards values that are optimal for producing the correct phonological output. Performance on inconsistent patterns such as -OWN is somewhat poorer because training on a word such as TOWN has a negative effect on the weights from the point of view of BLOWN and vice versa. In such cases, given sufficient training on the words, the model produces output that is closer to the correct pronunciation than to the alternative pronunciation of the inconsistent spelling pattern; however, error scores (the discrepancy between computed and veridical phonological codes) are larger than in the case of entirely consistent words.

The model's behaviour closely corresponds to that of human subjects asked to read words aloud; its performance is better on the words that subjects find easier and worse on the words they find more difficult. The accuracy of the model is such that it correctly simulates latency differences on the order of 15-25 msec. Many earlier studies of how subjects name different types of words aloud employed taxonomies of word types based on different assumptions about the nature of pronunciation rules and the perceptual units relevant to pronunciation (see Patterson and Coltheart 1987 for review). The model shows that the correct generalizations about naming performance derive from a deeper principle concerning the learning process.

These results have important theoretical implications. The inconsistency effects, which are exhibited by people and correctly simulated by the model, are not predicted by standard dual-route models, in which the fundamental distinction is between rule-governed words and exceptions. This dichotomy is not rich enough to capture facts about human performance. Highly regular words such as MUST and highly exceptional words such as CORPS represent different extremes on a continuum of spelling-sound consistency. Inconsistent words such as TINT or LEAF represent intermediate cases; they appear to be regular, 'rule-governed' items, but the naming of these items is in fact affected by knowledge of exception-word neighbours such as PINT and DEAF. Consistency effects are somewhat smaller than suggested by Glushko's (1979) original work but they can be detected with careful experimentation (see Jared et al. 1990), and they are theoretically important. Any number of theories can explain why a word with an irregular pronunciation might be more difficult to name than a regular word. However, the data indicate that differences among word types in terms of naming difficulty depend on the degree of consistency in the mapping between spelling and pronunciation. These differences in degree are realized in the model by the weights on connections, which reflect the aggregate effects of training on a large corpus of words.

Consistency of the past tense. With this background in hand, we can return now to the past tense. I previously argued that past tense inflection is analogous to spelling-sound correspondences in important respects and suggested that the two sets of phenomena might be explained by similar sorts of computational mechanisms. If this analysis is correct, it predicts that we should be able to observe analogous behavioural phenomena in the two domains. In particular, there should be consistency effects in the generation of the past tense. According to Pinker and Prince, WALK-WALKED is rule-governed and TAKE-TOOK is 'listed.' The 'Glushko question' for verbs, then, is what about BAKE-BAKED which, like TINT in the domain of spelling-sound correspondences, is rule-governed but inconsistent; its enemies are TAKE-TOOK and MAKE-MADE. Other examples are MIND-MINDED (inconsistent because of FIND-FOUND) and PUN-PUNNED (the enemy is RUN-RAN).[4]

Maggie Bruck and I (Seidenberg and Bruck 1990) examined these types of words in an on-line production task. On each trial subjects were shown a verb in the present tense, such as BAKE. Their task was either to name the word aloud or to generate its past tense. Subjects performed the tasks in two sessions separated by at least a week. Each subject performed both tasks; order of tasks was counterbalanced across subjects. The stimuli included fifty 'rule-governed' verbs with entirely regular past tenses, and fifty verbs with regular past tenses but one or more irregular neighbours (the 'inconsistent' items). There were also fifty-eight verbs with irregular past tenses included as filler items to keep subjects attending closely to the task. The regular and inconsistent items were closely matched in terms of properties of both the present and past tenses. We also chose the items so that, on average, the items in the two conditions had the same number of regular past tense neighbours. In this way we attempted to ensure that both the regular and inconsistent items involve 'rules' that are used about equally often in the language. The only systematic difference between the conditions was that the inconsistent items have enemies. The principal goal was to examine how past tense generation latencies relate to consistency. The present tense naming task was included in order to be certain that any differences in past tense generation times were not due to differences in the processing of the base words. The predictions should be clear: If the regular past tense is generated by rule, inconsistent past tenses such as BAKE-BAKED should yield the same results as entirely regular pasts such as WALK-WALKED. However, if the pattern of consistency across

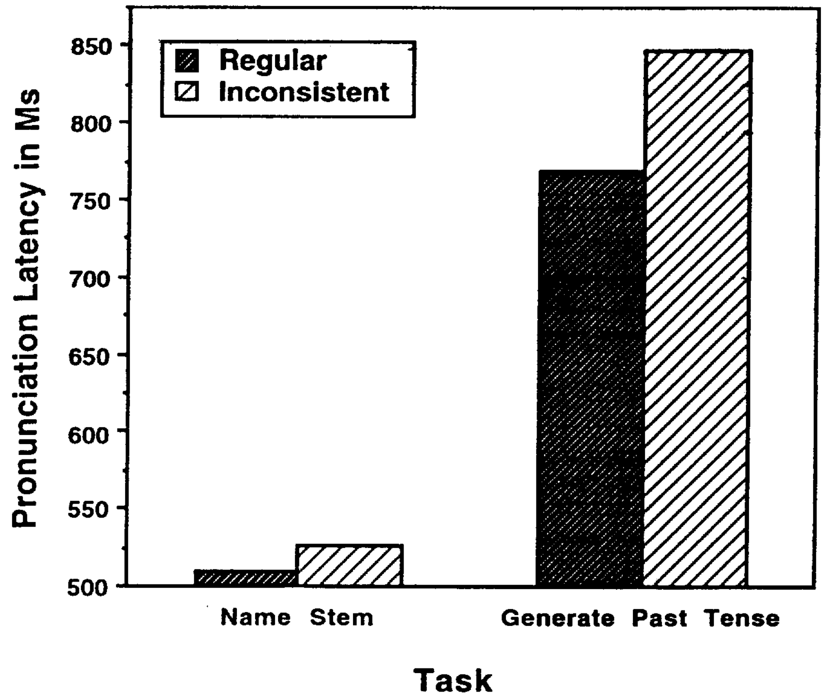

Figure 4.2: Results of the Seidenberg and Bruck (1990) study on past tense generation. Stimuli in both conditions were 'rule-governed'; however, the inconsistent items have enemies.

a neighbourhood of similarly-spelled forms is relevant, it should be harder to generate the past tenses of the inconsistent words.

The results, summarized in Figure 4.2, exhibit a strong, statistically reliable consistency effect. There were also 2 percent more errors in the inconsistent condition. In contrast, naming latencies for the present tense bases did not differ reliably in either latency or errors. One other interesting result was that for the inconsistent words, the latency to generate the past tense was related to the number of enemies, r (48) = .38, $p < .01$. The inconsistent words varied in terms of the number of enemies. Thus, items such as PICK-PICKED, whose only enemy is STICK-STUCK, were easier than items such as BLINK-BLINKED (whose enemies include SINK-SANK, THINK-THOUGHT, etc.).

The simplest interpretation of these results is that the regular past tense is not generated by rule. Rather, there is a computation over a neighbourhood of similarly-spelled patterns. As in the case of spelling-sound correspondences, it is the degree of consistency that captures the relevant generalization, not whether the item is 'rule-governed' or 'listed.'

As it happens, Prasada et al. (1990) also reported an experiment in which subjects generated the past tenses of verbs. They arrived at a somewhat different conclusion, however. Their experiment involved regular, rule-governed items and exceptions. In one case, they varied the frequencies of the present tense stems (as measured by the Francis and Kucera 1982 norms); thus, there were separate groups of high and low frequency regular and irregular items. The groups were equated, however, in terms of the frequencies of their past tenses. Hence, base word frequencies varied but past tense frequencies were the same. In a second condition, the stimuli in the two groups were equated in terms of base word frequencies but varied in terms of past tense frequencies (high versus low). As in our experiment, subjects saw the base word and generated the past tense. Prasada et al.'s results indicate that whereas the frequency of the base form affects the generation of both regular and irregular past tenses, the frequency of the past tense form itself was only relevant for irregular pasts. That is, the difference in frequency between TOOK (high frequency irregular) and BENT (low frequency irregular) affected response latencies, but the difference between LOOKED (high frequency regular) and BASKED (low frequency regular) did not.

Prasada et al. interpreted these results as support for the dual-route model. According to this account, regular past tenses are generated by rule. Overall latencies therefore consist of two components: the latency to identify the present tense stem (i.e., lexical access for LOOK) and a constant reflecting the time needed to apply the rule. Latencies to generate irregular past tenses also consist of two components: lexical access for the stem (for example, TAKE) and the amount of time it takes to find the irregular past tense listed in the mental lexicon. Importantly, the latter component is not a constant; it depends on the frequency of the word, under the assumption that the search process is frequency-ordered. It follows that the frequency of the base word affects both regular and irregular past tense generation, but the frequency of the past tense only affects the irregulars. Insofar as the data were in accord with these predictions, they were seen as confirming the dual-route account.

Bruck and I obtained very similar results using a slightly different design. The stimuli in our study were forty verbs with regular past tenses and forty with irregular past tenses. The present tense stems in the two conditions were equated in terms of Kucera and Francis (1967) frequency, length, and initial phoneme. Thirty subjects performed the two tasks described previously: naming the present tense base words aloud and generating the past tense forms. The tasks were again performed in separate sessions several days apart, with the order of tasks counterbalanced across subjects. The results are summarized in Figure 4.3. As in our previous experiment, naming latencies for the two types of present tense stems did not differ and they both yielded less than 1

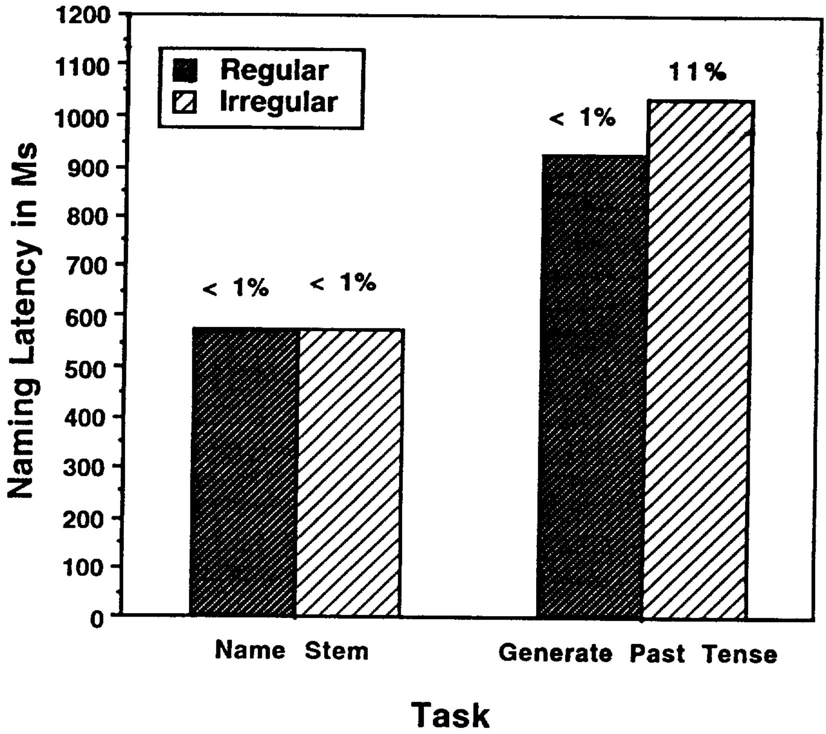

Figure 4.3: Results of the Seidenberg and Bruck (1990) study on past tense generation. Stimuli were either regular, rule-governed items or irregular, exception items.

percent errors. However, it took longer to generate the irregular past tense, and these words produced about 10 per cent more errors. Subtracting the stem naming latency from the past tense generation latency yields a net generation effect of 349 msec for the regular past and 456 msec for the irregular past. Thus, familiar irregular past tenses take about 100 msec longer to generate, even for skilled college student readers.

We also addressed the role of frequency by performing a median split on the lemma frequencies of the present tense stems. This yielded groups of high and low frequency stems for regular and irregular past tenses. Figure 4.4 presents the net past tense generation effects in these groups. Two findings should be noted: First, the difference between the regular and irregular conditions is larger for lower frequency words than high; second, there is a frequency effect for irregular past tenses but not for regular past tenses. This pattern is consistent with Prasada et al.'s results, which they took as evidence for the dual-route model.

Note, however, that the same pattern of results has repeatedly been observed in studies of spelling-sound correspondences. In these studies, the regularity effect (the difference in latency for exception

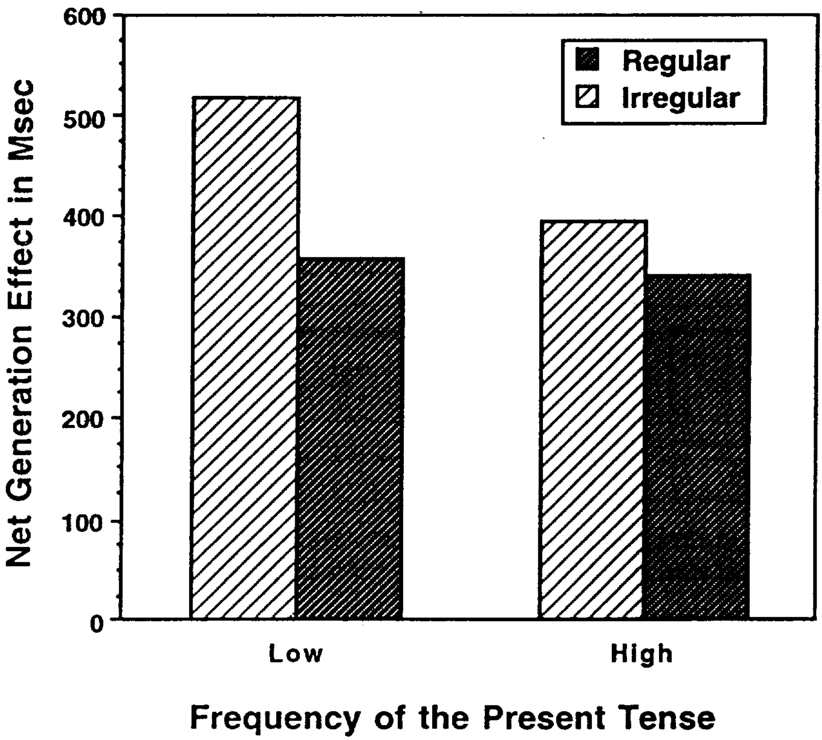

Figure 4.4: Data from the same experiment as in Figure 4.3, broken down by frequency.

words such as HAVE and regular words such as MUST) is larger for lower frequency words (Seidenberg 1985c). At the same time, frequency effects are smaller for regular words than they are for exceptions. The absolute size of the frequency effect for regular words depends on the range of frequencies sampled, but the effects are clearly smaller than for exceptions. Figure 4.5 presents the results of one representative study, by Waters and Seidenberg (1985), illustrating these effects. Data concerning the model's performance on the same words are also presented.

The Seidenberg and McClelland model simulates this frequency by regularity interaction quite closely. Hence, it exhibits the pattern of behaviour that Prasada et al. took as evidence for the dual-route model – even though it only has a single route. Specifically, it exhibits both the effect they interpreted as evidence for rule-use (minuscule frequency effects for regular items) and the effect that provided evidence for lexical-lookup (larger frequency effects for irregular items). However, it shows that these effects derive from the same source, namely, the effects of repeated changes to the weights during the training phase. Seidenberg and McClelland provided a detailed explanation of the factors that govern the model's behaviour. Performance on any given word is a

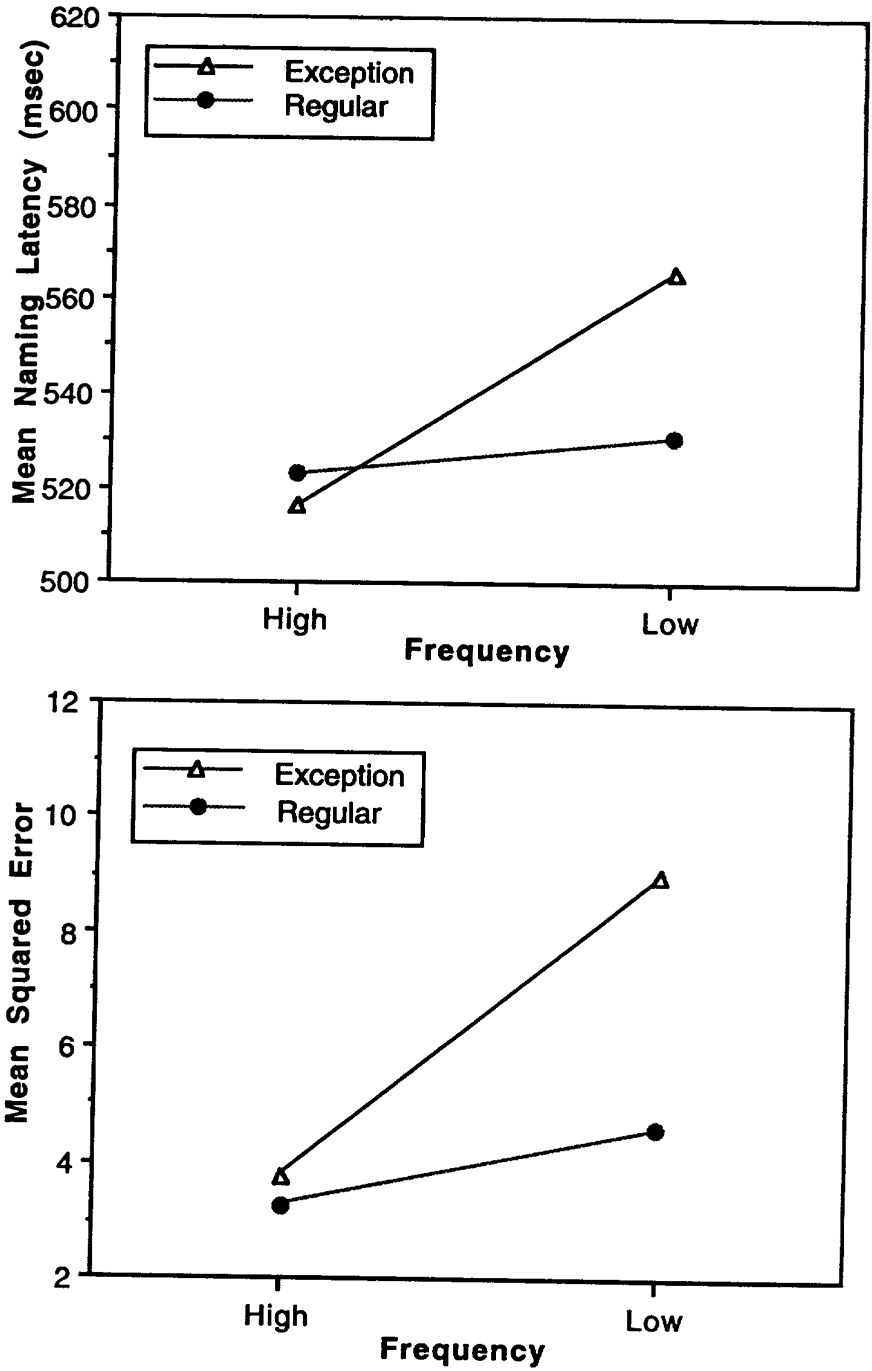

Figure 4.5: Data from an experiment by Waters and Seidenberg (1985) illustrating the frequency by regularity interaction in the naming of monosyllabic words.

function of the entire ensemble of training experiences. This is because all the changes to the weights that occur during learning are superimposed on each other. Hence, the weights reflect the aggregate effects of training on the entire corpus. For a given word, the factor that has the biggest impact on performance is the number of times the model was trained on the word itself. In this way lexical frequency has an impact on performance. However, performance is also affected by exposure to similarly-spelled neighbours; thus, performance on GAVE is affected by exposure to GAVE but also by neighbours such as SAVE and HAVE. There are also small effects due to more remote neighbours such as GIVE or MATE. The effects of these neighbours modulate effects of lexical frequency. As the number of neighbours (specifically, friends) goes up, the effects of number of exposures to the word itself decrease. Intuitively, mastering the pronunciation of GAVE is not highly dependent on exposure to GAVE because the model also benefits from SAVE, PAVE, and RAVE. In contrast, mastering the irregular pronunciation of HAVE is highly dependent on sufficient exposure to the word itself. Thus, frequency of exposure has a bigger impact on irregular patterns than on regular ones. For highly regular patterns with many friendly neighbours, effects of lexical frequency may be washed out entirely.

The claim here is that exactly the same factors govern latencies to generate the past tense. The regular, rule-governed patterns are highly productive. Hence, learning the past tense of words such as LOOK or LIKE is not highly dependent on the frequency of exposure to them. The correct past tenses can also be inferred on the basis of exposure to other regular, rule-governed forms. In contrast, learning the correct past tense of TAKE requires exposure to TOOK; therefore, performance is highly dependent on frequency. It follows from this view that frequency of the past tense should be a salient factor for irregular pasts but not for regular pasts, as in the data.

In summary, both Prasada et al. and Seidenberg and Bruck observed for past tense generation a frequency by regularity interaction like the one that has been observed for spelling-sound correspondences. The Seidenberg and McClelland model provides a simple account of the latter effect in terms of changes to the weights during learning. This single-process model obviates the need for a set of rules and a list of exceptions and suggests that a similar account should apply to past tense generation. Moreover, the Seidenberg and McClelland model also correctly predicts the effects of spelling-sound consistency observed in studies such as Glushko's (1979) and Jared et al.'s (1990), whereas the dual-route model did not. These consistency effects have now also been observed in past tense generation as well.

Error data. Subjects' errors on the past tense generation task also provide information that helps to differentiate between the theoretical alternatives. It is necessary to consider for a moment how errors might be generated within a dual-route model. The proposal is that regular past tenses are generated by rule rather than by being listed in the lexicon; irregular past tenses are produced by finding the forms in a memory list rather than by rule-application. A problem that arose with regard to the earlier, dual-route model of pronunciation was how the reader would know which route to use for any given input word. That is, if words are not labelled as 'regular' or 'irregular,' how does the reader know whether to pronounce by rule or to search for the irregular form? The usual answer to this is that both routes are tried in parallel, with a race between them (Meyer et al. 1974; Paap and Noel 1989). This proposal introduces other problems (for example, what happens when the routes yield different pronunciations), although I would say that they are unresolved rather than necessarily wholly intractable. In any case, Prasada et al. (1990) did not present a detailed process model addressing these issues. It would have to be assumed, however, that somehow the subject knows whether to apply a rule or to search the lexicon in generating the past tense for a given word. For example, words could carry tags indicating whether their past tenses are regular or irregular. An error would then result from reading the tag incorrectly; words with regular past tenses might be treated as irregular and vice versa. If a verb with an irregular past tense were mistakenly treated as though it were regular, the subject should produce a regularization error (for example, RUN-RUNNED). If a verb with a regular past tense were mistakenly treated as though it were irregular, however, it is not clear what kind of error should result. The subject would presumably initiate an unsucessful search through the mental lexicon. Under these conditions, the subject might respond with another irregular past tense found in the list, apply the regular rule as a default, or make no response at all.

Subjects' actual errors suggest a somewhat different picture. Errors seem to result from drawing incorrect analogies to neighbours. Table 1 presents the errors that occurred in generating past tenses for inconsistent verbs such as SIGHT in the first Seidenberg and Bruck (1990) experiment. Recall that the correct responses to these verbs always involve the regular, -ED pattern. Some of the errors (such as STREAK-STRUCK) are congruent with the dual-route model's suggestion that errors in generating the regular past tense would come about by mistakenly searching the list of irregular past tenses stored in memory and producing one of them as output. Thus, STRUCK is the correct irregular past tense of STRIKE not STREAK. However, most of the errors were not of this sort. The most frequent error was one in which the subject produced an incorrect past tense that was analogous to an irregular past tense in a nearby neighbourhood. Thus, SIGHT was pronounced SOUGHT by analogy to FIGHT-FOUGHT; THRIVE-THROVE by analogy to DIVE-DOVE, and GLIDE-GLID by analogy to HIDE-HID. Another interesting set of errors occurred when subjects incorrectly produced past tenses that were identical to the present tense forms. These are analogous to rare items such as HIT and BEAT, which have identical present and past tense forms. Thus, the subjects responded with BLIND as the past tense of BLIND (instead of BLINDED) and SKID as the past tense of SKID (instead of SKIDDED).

Errors such as GLIDE-GLID suggest that subjects were generating the past tenses by analogy to other forms. GLIDE sounds like HIDE,

Table 4.1: Errors on the Regular But Inconsistent Verbs Used in Seidenberg and Bruck's Experiment 1

Vowel changes:	
sight-sought	(13)
thrive-throve	(5)
weave-wave	(3)
glide-glid	(2)
squeeze-squoze	(2)
streak-struck	(2)
streak-stroke	(1)
No change errors:	
blind	(2)
brand	(1)
free	(2)
skid	(1)
slow	(2)
thread	(2)
wing	(1)
Other errors:	
sight-saw	(4)
lean-leant	(4)
smell-smelt	(7)
trust-thrusted	(1)
rig-ringed	(1)
streak-shrieked	(1)
No response:	(7)

Note: Number of errors given in parentheses.

therefore the past tense should sound like HID. The same process would also produce errors such as SIGHT-SOUGHT and 'no change' errors as well. Thus, the errors that occurred for the regular but inconsistent verbs seem to reflect the effects of similarly spelled or pronounced neighbours. This is congruent with the hypothesis that the past tense is generated by means of a computation that reflects relationships among a neighbourhood of words rather than by simply applying a rule. The 'analogy' process is realized in the weights, which reflect the degree of consistency in the mapping between input (for example, present tense) and output (for example, past tense) forms.

For verbs with irregular past tenses, the only error that is expected on the dual-route account would be a regularization such as RUN-RUNNED. Some of these errors did occur (Table 4.2): KNOWED, HURTED, and FIGHTED are examples. However, subjects also produced a variety of other errors. Some were no-change errors; thus, in the case where the subject generated KNOW as the past tense of KNOW, he produced an irregular, no-change past tense but not one that would be 'listed' in the lexicon as the past tense of some other verb. Subjects also produced analogies that formed nonwords, such as SEEK-SOOK. These errors are especially important because they could not result from either mistakenly applying a rule or accessing the incorrect entry in the mental lexicon. Subjects also incorrectly produced some past participles such as SEE-SEEN. Again, the errors seem to reflect relationships between the stimulus verb and similarly spelled or pronounced neighbouring words rather than mere application of a rule. The very similar types of errors produced for both regular and irregular past tenses strongly suggest that these forms are generated by means of a common process. I should add that subjects produce similar errors in studies of naming monosyllabic words aloud. In reading familiar irregular words such as HAVE or DEAF, subjects sometimes produce regularizations (/hAv/, /dEf/), but they also produce other types of errors (HAVE-/hIv, DEAF-/det/). Conversely, regular but inconsistent words are sometimes incorrectly read by analogy to exceptions, for example, GAVE-/gav/ or TOWN-/tOn/.

In summary, three phenomena have been observed in these studies of past tense generation. First, subjects exhibit consistency effects for entirely regular, rule-governed past tenses such as BAKE-BAKED; second, frequency effects are bigger for the irregular past tense than for the regular past tense; third, errors reflect relationships between a word and its neighbours. These effects also occur in the domain of spelling-sound correspondences, and the Seidenberg and McClelland model simulates them closely. These observations strongly suggest that it

Table 4.2: Errors on the Irregular Verbs Used in Seidenberg and Bruck's Experiment 2

Regularizations	
knowed	(1)
hurted	(2)
kneeled	(2)
fleed	(3)
holded	(1)
fighted	(1)
ridded	(3)
No change errors:	
know	(1)
fly	(1)
flee	(3)
fall	(1)
run	(1)
draw	(1)
slide	(1)
stand	(1)
mean	(4)
stick	(1)
Past tenses of other words:	
flee-flew	(8)
rid-rode	(3)
run-rung	(1)
Past participles:	
seen	(1)
given	(1)
ridden	(3)
eaten	(2)
driven	(2)
broken	(1)
frozen	(1)
beaten	(2)
stolen	(1)
Other errors:	
seek-sook	(1)
dig-dag	(1)
strike-stroke	(2)
lose-loose	(1)
swim-swum	(2)
fall-fail	(1)
mean-meent	(1)
give-gaved	(1)
steal-stoled	(2)

Note: Number of errors given in parentheses.

would be worth pursuing the single-route, connectionist approach to modelling the past tense.

Generalization. Before closing this discussion the important issue of generalization must be considered. Perhaps the quintessential property of a rule is that it can account for both known instances and the ability to generalize to novel instances. Thus, in the classic Berko (1958) study of children's knowledge of morphology, the inference that they had formulated a rule for the plural was based on the production of novel forms such as WUGS. It is obvious that people can generalize the past tense rule as well; what is the past tense of GLORP if not GLORPED? The single-route, connectionist interpretation of such behaviour is that novel forms are produced by the same mechanism as known forms, namely, the net. Thus, a model might be trained on a variety of verbs, resulting in changes to the weights that reflect facts about the past tense. The same weights would then be used in generating a novel past tense such as GLORPED. One of the important tests for future models of the past tense will be to determine if they generalize in appropriate ways. I myself doubt whether this will be a serious problem, assuming the model is trained in a way that faithfully reflects facts about the distribution of regular and irregular past tenses in the language. The system is overwhelmingly regular; and the weights will come to reflect this fact, making it likely that the regular past tense will be attached to almost any novel input. The only exceptions would be cases where the novel input happens to fall within one of the clusters of irregular pasts, for example, TING might be given the past tense TANG on the basis of neighbours such as RING-RANG and SING-SANG. There is some evidence that children produce such forms (Bybee and Moder 1983). I think it is likely to be more of a challenge to get a model to correctly produce irregular past tenses given the overwhelming degree of regularity in the system.

Although the Seidenberg and McClelland (1989) model illustrates how generalization occurs within a simple network, it also raises questions as to whether such networks can achieve performance that is as good as peoples'. The model was trained on monosyllabic words, and correctly generalizes when presented with simple nonwords such as NUST or RIKE. Glushko (1979) had observed that nonword naming latencies also exhibit consistency effects; thus nonwords such as NUST, from the entirely consistent -UST neighbourhood, are named more rapidly than are nonwords such as MAVE, from the inconsistent -AVE neighbourhood. Although the standard view is that novel forms are generated by applying the regular rules, these results suggest that generation involves a network that encodes similarity and consistency relationships among pools of neighbours, as in the case of words. The

Seidenberg and McClelland model also produces consistency effects for simple nonwords. The model does not perform as well as people in naming nonwords, however; it produces a larger percentage of errors on more difficult items such as JINJE or KEAD (Besner et al. 1990). As Seidenberg and McClelland (1990) noted, however, this behaviour of the model is closely related to the fact that whereas peoples' vocabularies are on the order of tens of thousands of words, the model was trained on only 2897. Thus, the model does well on nonwords that resemble items in the training corpus (for example, NUST) but poorly on unusual items such as JINJE. One way to view these results is that the model performs about as well as one might expect of a person who only knows 2900 monosyllabic words. Other aspects of the implementation, particularly the phonological representation that was used, also limit the model's performance on nonwords (see Seidenberg and McClelland 1990 for discussion). These observations suggest that the limitations that have been observed may not be insuperable. Still, given the concerns that have been raised regarding the capacities of simple backprop nets to generalize (McCloskey and Cohen 1989), it will be important to investigate this issue further – and considerable caution is in order. In the case of the past tense of verbs, it will be important to determine whether a network can both generate correct past tenses for known verbs and generalize even in the case of odd nonwords such as XPLK; even though XPLK has no close neighbours (or perhaps because of it?), we can agree that its past tense must be XPLKED. As I have noted, I think this is likely to be a tractable problem given the extreme regularity of the system but, this is a critical empirical question that must be addressed.

Conclusion

One of the most important contributions of the Pinker and Prince (1988) paper is that it provided a description of a broad range of phenomena that any adequate theory of the past tense must explain. I have suggested that it is by no means obvious that the phenomena they highlighted lie outside the scope of simple connectionist models, the failures of the McClelland and Rumelhart (1986) model notwithstanding.[5] To their list of phenomena should be added those that were uncovered in the Seidenberg and Bruck (1990) and Prasada et al. (1990) studies. These data are strongly reminiscent of that which were observed earlier in connection with spelling-sound correspondences.

If I am correct in suggesting that phenomena such as the past tense in English can be accommodated by a simple connectionist architecture, this would suggest a somewhat different relationship between linguistic theory and connectionist modelling than that implied by Pinker and Prince's 'implementational' view. The Seidenberg and McClelland (1989) model does not merely 'implement' the dual-route model of naming. Our analyses of the model indicate that it cannot be decomposed into components corresponding to a set of rules and a list of exceptions; thus, the two mechanisms of the dual-route model are not directly implemented. Moreover, the model behaves in ways that would not be predicted on the simple rules-and-exceptions view. The ways in which frequency and consistency of spelling-sound correspondences affect processing follow from an understanding of how learning works in a simple network employing distributed representations, not from the properties of the rule and lookup mechanisms proposed in earlier theories. It is a strong argument in the model's favour that people exhibit analogous behaviours. In the case of the past tense, there is as yet no implemented model that addresses all (or even most) of the relevant phenomena, but thinking about the generation of past tense forms in connectionist terms has already led to predictions that have been confirmed in behavioural studies. I suggest that these phenomena would not have been discovered without an understanding of how learning works in connectionist networks. To the extent that the connectionist framework both accounts for various facts and generates novel, correct predictions, it cannot be said to be simply 'implementing' the rules.

Is it the case that phenomena of the sort I have described (for example, consistency effects) are wholly incompatible with the rules-and-exceptions approach? Certainly not. One response to data of this sort would be to modify the dual-route model in order to accommodate them. That is what happened in the case of spelling-sound correspondences, and it could as well occur with respect to verbs. Thus, Patterson and Coltehart (1987) describe various modifications of the dual-route model intended to cope with the Glushko consistency effects. Similarly, Pinker (1991) has recently described a modified dual-route model in which some forms are generated by means of an associative net, and others are generated by rule. In response to the Seidenberg and Bruck (1990) results, Pinker now suggests that the associative net applies to inconsistent words, such as BAKE or MIND, previously thought to fall within the purview of the rule-component. Whether or not this move will be successful is unclear. It certainly introduces some important questions as to how such a system would ever be learned. The child would have to learn that BAKE-BAKED and TAKE-TOOK are processed by one mechanism, even though they are superficially quite different, whereas BAKE-BAKED and LIKE-LIKED are processed by separate mechanisms, even though they are superficially quite similar.

Still, it is not inconceivable that such problems could be resolved. Certainly, in the absence of an implemented connectionist model of the past tense that is at least descriptively adequate, Pinker's alternative is equally viable.

Note, however, that this discussion involves a very different relationship between connectionism and linguistic theory than Pinker and Prince (1988) envisioned. Their idea was that connectionist models could only 'implement' the types of knowledge structures and processing mechanisms uncovered by linguistic theory. That would involve investing the connectionist network with properties independently established within linguistic theory. What Pinker (1991) seems to have in mind is the exact opposite: investing a rule-based system with properties independently discovered on the basis of connectionist modelling. If it turns out that rules apply to words from consistent neighbours but not to words from inconsistent neighbourhoods, change the notion of rule so that it obeys this principle. If it turns out that facility in producing the past tense depends on how often a pattern occurs in the language, assign frequencies to the rules. In general, there is sufficient elasticity in the notion of 'rule' to permit a rule-based account to accommodate nearly any pattern of data. Unless there are constraints on the properties of rules, there would not seem to be any limits to their range of applicability.

What is important, of course, is not whether one theory can mimic another but, rather, from where the correct generalizations are derived. In this regard, the approach that I have advocated is profoundly different from Pinker and Prince's. I began this paper by suggesting that connectionist models can be properly understood with reference to the learnability notion that has been central to studies of language acquisition. The models I have been describing are systems that learn under certain specifiable constraints. From this point of view, it is critical to understand such things as the initial state of the system, the input to the system, and the capacity of the system to learn. The generalizations that govern performance derive from the interaction of these factors. It is because they are so critical to a model's performance that Pinker and Prince's criticisms of the Rumelhart and McClelland (1986) model (for example, of the training regime and the phonological representation) were so telling. The central, defining feature of this approach is that it is centred on the question of how a particular task is mastered (for example, learning a language). In a much simpler fashion, our models are also task-centred: they ask how a system can come to perform a task such as pronouncing words or generating the past tense.

Pinker and Prince's approach starts with a characterization of the knowledge of the adult – competence – and asks how this is achieved.

Thus, they assume a theory of inflectional morphology that distinguishes between rules and exceptions. This theory is primarily derived from distributional analyses of adult utterances rather than from the analysis of a task. The rules and lists of exceptions are attempts to rationalize the regularities implicit in this large set of observational data. As I have noted, the view that knowledge consists of rules and lists of exceptions is by no means universal among theoretical linguists. I would argue that this view is plausible only if one considers relatively simple systems such as inflectional morphology in English, and even in this simple domain there are phenomena suggesting that it is an oversimplification. Nonetheless, when Pinker and Prince turn to considering the acquisition process, they assume that a proper theory will necessarily respect the rule/list distinction. They are able to amass a large amount of data that are consistent with the distinction and attempt to sketch learning mechanisms that are compatible with it. Among the other approaches afforded by theoretical linguistics, however, is learnability. Ideally, what has to be independently motivated are the initial state of the system, the way in which knowledge is represented (for example, in terms of weights on connections), the input to the child, and the way in which learning occurs. From the interaction of these factors a certain type of competence necessarily follows. Thus, knowledge representations develop in the service of mastering a task. This contrasts with theories in which knowledge representations reflect generalizations derived from comparisons across adult utterances or across languages.

According to Rumelhart et al. (1986), the rules and the list of exceptions can be taken as simply an imprecise, higher level characterization of the behaviour of a complex system. Abstracting from the details of the Seidenberg and McClelland model, for example, one could say that it behaves as though it had induced the rules governing most words in the language but also represented the exceptions and not be wholly inaccurate. With the model in hand, however, one can see that it does not, in fact, implement anything like the rules or the list of exceptions previously envisioned. In fact, its behaviour departs systematically from what was expected on the rules-and-exceptions view. Thus, the fact that the model's behaviour can be summarized in a certain way should not blind us to how it actually works. And the virtue of having an implemented simulation model is that one can actually see.

Of course, there are very few models that achieve any kind of descriptive adequacy; even our model is severely limited in scope. It is absurdly ambitious to attempt to develop systems that mimic human behaviour in detail, and certainly very little has been achieved in this regard so far. Moreover, the difficulty of this task – and the limitations of scope that this imposes – ensures that any given model will simply be false insofar as it fails to be faithful to all of the phenomena of interest. These observations certainly justify Pinker and Prince's robust skepticism and suggest that theories of the sort they described will continue to play an important role and, in many domains, will continue to be the best accounts that can be achieved.

In my view, the great divide is not between linguistic theory and connectionism; it is between theories that are centred on the learnability question and those that treat it as secondary to characterizations of adult competence. I myself do not believe that the non-learnability approaches that are common in many areas of theoretical linguistics (for example, morphology) are powerful enough to converge on the correct characterizations of linguistic knowledge. Learnability questions are often acknowledged but in some domains they do not play a central role in theory development. I see connectionism as contributing in an essential way to achieving explanatory theories of a sort to which many linguists aspire.

Notes

1 I am quoting here from their description of a course they jointly offered at the Summer Institute of the Linguistic Society of America (Santa Cruz 1991).

2 I am quoting here from the title of a talk, 'Rules and Associations,' given at several locations. I heard it at a meeting convened by the McDonnell-Pew Foundation in San Diego 1990.

3 I include here the 'principles and parameters' approach to acquisition, even though it differs somewhat from the earlier learnability work.

4 Pinker and Prince (1989) appear to have sensed that the inconsistent items would behave differently than would entirely regular words: 'In contrast [to irregular verbs], regular verbs, *unless they are similar to an irregular cluster,* have no gradient of acceptability based on their phonology' (188; italics added).

5 Though I would not want the 'simple models' to be limited to feedforward nets trained using backprop.

References

Anderson, S.R. (1982). Where's morphology? *Linguistic Inquiry* 13:571-612

Aronoff, M. (1976). *Word Formation in Generative Grammar.* Cambridge, MA: MIT Press

Berg, G. (1991). Learning recursive phrase structure: Combining the strengths of PDP and X-bar syntax. TR 91-5, Department of Computer Science, SUNY Albany

Berko, J. (1958). The child's learning of English morphology. *Word* 14:150-77

Besner, D., Twilley, L., McCann, R., and Seergobin, K. (1990). On the connection between connectionism and data: Are a few words necessary? *Psychological Review* 97:432-46

Bromberger, S., and Halle, M. (1989). Why phonology is different. *Linguistic Inquiry* 20:51-70

Bybee, J.L. (1985). *Morphology: A Study of the Relation Between Meaning and Form.* Amsterdam/Philadelphia: John Benjamins

– and Moder, C. (1983). Morphological classes as natural categories. *Language* 59:251-70

Carr, T.H., and Pollatsek, A. (1985). Recognition of printed words: A look at current models. In D. Besner, T.G. Waller, and G.E. MacKinnon (eds.), *Reading Research: Advances in Theory and Practice* 5. New York: Academic Press

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

– (1981). *Lectures on Government and Binding.* Dordrecht, The Netherlands: Foris Publications

Coltheart, M. (1978). Lexical access in simple reading tasks. In G. Underwood (ed.), *Strategies of Information Processing.* New York: Academic Press

– (1987). Introduction. In M. Coltheart, G. Sartori, and R. Job (eds.), *Cognitive Neuropsychology of Language.* London: Erlbaum

Cottrell, G., and Plunkett, K. (1991). Using a recurrent net to learn the past tense. In *Proceedings of the Thirteenth Annual Meeting of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum

Elman, J. (1990). Finding structure in time. *Cognitive Science* 14:179-211

Fernald, A. (1984). The perceptual and affective salience of mothers' speech to infants. In L. Feagans, C. Garvey, and R. Golinkoff (eds.), *The Origins and Growth of Communication.* New Brunswick, N.J.: Ablex

Francis, W. N. and Kucera, H. (1982). *Frequency Analysis of English Usage: Lexicon and Grammar.* Boston: Houghton-Mifflin

Gluck, M., and Bower, G. (1988). Evaluating an adaptive network model of human learning. *Journal of Memory and Language* 27:166-95

Glushko, R.(1979). The organization and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance* 5:674-91

Halle, M. (1973). Prolegomena to a theory of word-formation. *Linguistic Inquiry* 4:3-16

– and Mohanan, K.P. (1985). Segmental phonology and modern English. *Linguistic Inquiry* 16:57-116

Hanna, P.R., Hanna, J.S., Hodges, R.E., and Rudorf, E.H. (1966). *Phoneme-Grapheme Correspondences as Cues to Spelling Improvement.* Washington: US Department of Health, Education, and Welfare

Hansen, D., and Rodgers, T.S. (1968). An exploration of psycholinguistic units in initial reading. In K.S. Goodman (ed.), *The Psycholinguistic Nature of the Reading Process.* Detroit: Wayne State University Press

Hare, M. (1990). The role of similarity in Hungarian vowel harmony: A connectionist account. *Connection Science* 2:123-50

Hinton, G.E. (1986). Learning distributed representations of concepts. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society* 1-12. Hillsdale, NJ: Erlbaum

Hirsh-Pasek, K., Treiman, R., and Schneiderman, M. (1984). Brown and Hanlon revised: Mothers' sensitivity to ungrammatical forms. *Journal of Child Language* 11:81-88

Hynd, G., and Semrud-Clikeman, M. (1989). Dyslexia and brain morphology. *Psychological Bulletin* 106:447-82

Jared, D., McRae, K., and Seidenberg, M.S. (1990). The basis of consistency effects in word naming. *Journal of Memory and Language*, 29:687-715

Kiparsky, P. (1982). Lexical morphology and phonology. In I.-S. Yang (Ed.), *Linguistics in the Morning Calm*. Seoul: Hanshin, 3-91

Kucera, H., and Francis, W.N. (1967). *Computational Analysis of Present-day American English*. Providence, RI: Brown University Press

Liberman, A.M., and Mattingly, A.G. (1985). The motor theory of speech perception revisited. *Cognition* 21:1-36

MacWhinney, B. (in press). Paper to appear in *Cognition*

McCloskey, M., and Cohen, N.J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G.H. Bower (ed.), *The Psychology of Learning and Motivation*, 23. New York: Academic Press

McRae, K., Jared, D., and Seidenberg, M.S. (1990). On the roles of frequency and lexical access in word naming. *Journal of Memory and Language* 29:43-65

Meyer, D.E., Schvaneveldt, R.W., and Ruddy, M.G. (1974). Functions of graphemic and phonemic codes in visual word recognition. *Memory and Cognition* 2:309-21

Newport, E., Gleitman, L., and Gleitman, H. (1977). Mother, I'd rather do it myself: Some effects and non-effects of maternal speech style. In C.E. Snow and C.A. Ferguson (eds.), *Talking to Children: Language Input and Acquisition*. New York: Cambridge University Press

Paap, K., and Noel, R.W. (1989) Dual-route models of print to sound: Still a good horse race. Paper presented at the annual meeting of the Psychonomic Society

Patterson, K. E., Marshall, J. C., and Coltheart, M. (1985). *Surface Dyslexia: Neuropsychological and Cognitive Studies of Phonological Reading*. London: Erlbaum

– and Coltheart, V. (1987). Phonological processes in reading: A tutorial review. In M. Coltheart (ed.), *Attention and Performance XII: Reading*. London: Erlbaum

Pinker, S. (1991). *Rules of Language*. Unpublished ms., MIT

– and Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* 28, 73-194

– and Prince, A. (1989). Rules and connections in human language. In R.G.M. Morris (ed.), *Parallel Distributed Processing: Implications for Psychology and Neurobiology*. Oxford, UK: Oxford University Press

Plunkett, K., and Marchman, V. (1991). U-shaped learning and frequency effects in a multilayered perceptron: Implications for child language acquisition. *Cognition* 38:43-102

Prasada, S., Pinker, S., and Snyder, W. (1990). Some evidence that irregular forms are retrieved from memory but regular forms are rule generated. Paper presented at the thirty-first annual meeting of the Psychonomic Society (November: New Orleans)

Rumelhart, D., Hinton, G., and Williams, R. (1986). Learning internal representations by error propagation. In Rumelhart, D.E., and McClelland, J.L., eds. (1986a), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* 1. Cambridge, MA: MIT Press

– and McClelland, J.L. (1986). On learning the past tenses of English verbs. In J. McClelland, and D. Rumelhart (eds.) *Parallel Distributed Processing* 2. Cambridge, MA: MIT Press

Seidenberg, M.S. (1985a). Evidence from great apes concerning the biological bases of language. In A. Marras and W. Demopoulos (Eds.), *Language Learnability and Concept Acquisition*. Norwood, NJ: Ablex

– (1985b). The time course of information activation and utilization in visual word recognition. In D. Besner, T. Waller, and G.E. MacKinnon (eds.), *Reading Research: Advances in Theory and Practice* 5. New York: Academic Press

– (1985c). The time course of phonological code activation in two writing systems. *Cognition* 19:1-30

– (1988). Cognitive neuropsychology and language: The state of the art. *Cognitive Neuropsychology* 5:403-26

– (1989). Reading complex words. In G. Carlson and M.K. Tanenhaus (eds.), *Linguistic Structure in Language Processing*. Dordrecht: Kluwer Academic Publishers

– and Bruck, M. (1990). Consistency effects in the generation of past tense morphology. Paper presented at the thirty-first meeting of the Psychonomic Society (November: New Orleans)

– and McClelland, J.L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review* 96, 523-68

– and McClelland, J.L. (1990). More words but still no lexicon: Reply to Besner et al. *Psychological Review* 97: 447-52

– and Petitto, L.A. (1987). Communication, symbolic communication, and language. *Journal of Experimental Psychology: General* 116:279-87

Sejnowski, T.J., and Rosenberg, C.R. (1986). NETtalk: A parallel network that learns to read aloud. Baltimore, MD: Johns Hopkins University EE and CS Technical Report JHU/EECS-86/01

Selkirk, E. (1982). *The Syntax of Words*. Cambridge, MA: MIT Press

– (1984). *Phonology and Syntax: The Relation Between Sound and Structure*. Cambridge, MA: MIT Press

Spencer, A. (1991). *Morphological Theory*. Oxford, UK: Blackwell

Spoehr, K.T., and Smith, E.E. (1973). The role of syllables in perceptual processing. *Cognitive Psychology* 5:71-89

Stabler, E. (1991). Book to be published by MIT press

Venezky, R. (1970). *The Structure of English Orthography.* The Hague: Mouton

Waters, G.S., and Seidenberg, M.S. (1985). Spelling-sound effects in reading: Time course and decision criteria. *Memory and Cognition* 13:557-72

Wijk, A. (1966). *Rules of Pronunciation for the English Language.* Oxford University Press