Running Head: CONNECTIONIST PSYCHOLINGUISTICS

Connectionist psycholinguistics: The very idea

Morten H. Christiansen

Southern Illinois University, Carbondale

Nick Chater

University of Warwick

Abstract

Connectionist psycholinguistics: The very idea

Introduction

What is the significance of connectionist models of language processing? Will connectionism ultimately replace, complement or simply implement the symbolic approach to language? Early connectionist models attempted to address this issue by showing that connectionist models could, in principle, capture aspects of language processing and linguistic structure. Little attention was generally paid to the modeling of data from psycholinguistic experiments. However, we suggest that connectionist language processing has matured and that the field is now moving forward into a new phase in which closer attention is paid to detailed psycholinguistic data. This book provides the first comprehensive overview of work within the emergent field of *"connectionist psycholinguistics"*—connectionist models that make close contact with psycholinguistic results.

But how are we to assess the models within this emerging new area of research? We suggest that computational models of psycholinguistic processing, whether connectionist or symbolic, should attempt to fulfill three criteria: a) *data contact*, b) *task veridicality*, and c) *input representativeness* (Christiansen & Chater, 2000). Data contact refers to the degree to which a model provides a fit with psycholinguistic data. We distinguish here between primary and secondary data contact. Primary data contact involves fitting results from specific psycholinguistic experiments (e.g., reaction time data), whereas secondary data contact involves fitting general patterns of behavior (e.g., experimentally attested developmental changes in language processing) rather than specific results. Task veridicality refers to the degree of match between the task facing people and the task given to the model. Although a precise match is typically difficult to obtain, it is important to minimize the discrepancy. For example, much early work on modeling the English past tense suffers from low task veridicality (e.g., Rumelhart & McClelland, 1986—but see e.g., Hoeffner, 1997, for an exception) because models are trained to map verb stems to past tense forms, a task unlikely to be relevant to children's language acquisition. Input representativeness refers to the degree to which the information given to the model reflects what is available to a person or child. For example, the computational modeling of morphology suffers from the lack of good training corpora of high input representativeness with which to train the models. This problem is most serious for non-English morphology, making it problematic to make a priori conclusions about the feasibility of connectionist accounts in the area (e.g., Berent, Pinker & Shimron, 1999).

It is important also to take stock of where symbolic models stand on our three criteria for computational psycholinguistics. Interestingly, few symbolic models make direct contact with psycholinguistic data. Most of the exceptions are within the study of sentence processing where some comprehensive models of word-by-word reading times exist (e.g., Gibson, 1998; Just & Carpenter, 1992), and have a reasonable degree of task veridicality. More generally, however, symbolic models appear to pay little attention to task veridicality. Indeed, the rule-based models of the English past tense (e.g., Pinker, 1991) involve the same stem-to-past-tense mappings as the early connectionist models, and thus suffer from the same low task veridicality. Input representativeness is often ignored in symbolic models, in part because learning plays a minimal role in the performance of these models, and in part because symbolic models tend to be focused on more abstract fragments of language, rather than the more realistic language input that some connectionist models can handle. Low input representativeness may, for these reasons, actually inflate performance for many types of symbolic models, whereas the opposite tends to be true of connectionist models.

Currently, then, connectionism appears to provide a better framework for detailed psycholinguistic modeling than does the symbolic approach. For many connectionists, the advantages of this framework for doing computational psycholinguistics derive from a number of properties of the connectionist models:

Learning. Connectionist networks typically learn from experience, rather than being fully prespecified by a designer. By contrast, symbolic computational systems, including those concerned with language processing, are typically, but not always, fully specified by the designer.

Generalization. Few aspects of language are simple enough to be learnable by rote. The ability of networks to generalize to cases on which they have not been trained is thus a critical test for many connectionist models.

Representation. Because they are able to learn, the internal codes used by connectionist networks need not be fully specified by a designer, but are devised by the network so as to be appropriate for the task. Developing methods for understanding the codes that the network develops is an important strand of connectionist research. While internal codes may be learned, the inputs and outputs to a network generally use a code specified by the designer. These codes can be crucial in determining network performance. How these codes relate to standard symbolic representations of language in linguistics is a major point of contention.

Rules vs. *Exceptions.* Many aspects of language can be described in terms of what have been termed "quasi-regularities"—regularities which are usually true, but which admit some exceptions. According to the symbolic descriptions used by modern linguistics, these quasi-regularities may be captured in terms of a set of symbolic rules, and sets of exceptions to those rules. Symbolic models often incorporate this distinction by having separate mechanisms which deal with rule-governed and exceptional cases. It has been argued that connectionist models provide a single mechanism which can pick up general rules, while learning the exceptions to those rules. While this issue has been a major point of controversy surrounding connectionist models, it is important to note that attempting to provide single mechanisms for rules and exceptions is not essential to the connectionist approach; one or both separate mechanisms for rules and exceptions could themselves be modeled in connectionist terms (Coltheart, Curtis, Atkins & Haller, 1993; Pinker, 1991; Pinker & Prince, 1988). A further question is whether networks really learn rules at all, or whether they simply approximate rule-like behavior. Opinions differ concerning whether the latter is an important positive proposal, which may lead to a revision of the role of rules in linguistics (Rumelhart & McClelland, 1986; see also Smolensky, 1988), or whether it is a fatal problem with connectionist models of language processing (Marcus, 1998; Pinker & Prince, 1988).

These four properties all play an important role in the models described in Part I of this volume as well as in the appraisals of connectionist psycholinguistics presented in Part II.

Organization of this Volume

Part I of this volume, *The State of the Art*, presents the current state of the art in connectionist psycholinguistics with specific models from five key areas: speech processing, morphology, sentence processing, language production, and reading aloud. Part II, *Future Prospects*, then provides three contrasting perspectives on the field, from leading researchers working on computational models of human natural language processing. In this section, we briefly outline the contents of each chapter in turn.

Part I begins with Chapter 2, *Connectionist psycholinguistics in perspective*, by Morten Christiansen and Nick Chater. This chapter provides an in-depth introduction to the field of connectionist psycholinguistics, and sets the context for the rest of the volume. The historical roots of the connectionist approach to language processing are traced and key themes that arise throughout different areas of connectionist psycholinguistics are highlighted. The chapter also provides a detailed review of the five key empirical areas described in the chapters comprising the rest of this part of the book, highlighting the interplay between connectionist modeling and empirical research. This review indicates that connectionist psycholinguistics has already had a significant impact on the psychology of language, and suggests that connectionist models are likely to have an important influence on future research. With this review in place, the subsequent chapters in this part of the book present central recent developments in the field.

Chapter 3, *Simulating parallel activation in spoken word recognition* by Gareth Gaskell and William Marslen-Wilson, concerns the connectionist modeling of speech processing. A critical property of the perception of spoken words is the transient ambiguity of the speech signal. Speech information is spread out across time, and early on in the processing of a word, the speech information will be compatible with more than one lexical item. In localist models of speech perception this property is captured by allowing the parallel activation of multiple independent lexical representations. Gaskell and Marslen-Wilson examine how this property can be accommodated in a distributed model of speech perception, in which word representations are not independent. In this case, an approximation to the activation of more than one representation is possible by activating a "blend" of the different distributed representations. Statistical analyses of vector spaces show that coactivation of multiple distributed representations is inherently noisy, and depends on parameters such as sparseness and dimensionality. Furthermore, the characteristics of coactivation vary considerably, depending on the organization of distributed representations within the mental lexicon. This view of lexical access is supported by analyses of phonological and semantic word representations, which provide an explanation of a recent set of experiments on coactivation in speech perception (Gaskell & Marslen-Wilson, 1999). More generally, this work illustrates a tight interplay between connectionist psycholinguistic modeling and experimental psycholinguistic research. Thus, the model provides for a good primary data contact and reasonable input representativeness, but suffers from a relatively poor task veridicality because of the abstract nature of the simulations.

Chapter 4, *A connectionist model of English past tense and plural morphology*, by Kim Plunkett and Patrick Juola concerns what has been one of the most controversial domains to which connectionist research has been applied: morphological processing. Theorists advocating a symbolic perspective have frequently taken morphology as a paradigmatic case of a “rule + exception” mapping. A rigid symbolic rule, which specifies a regular morphological mapping, is presumed to be supplemented with a set of explicit exceptions, which are assumed to be processed by a very different mechanism. In line with much connectionist work in this area, Plunkett and Juola make the opposite assumption—that a single mechanism explains both “rule” and “exception” cases in morphological processing. Specifically, they model the acquisition of English noun and verb morphology using a single connectionist network. The network is trained to produce the plurals and past tense forms for a large corpus of monosyllabic English nouns and verbs. The developmental trajectory of network performance is analyzed in detail and is shown to mimic a number of important features of the acquisition of English noun and verb morphology in young children. These include an initial error-free period of performance on both nouns and verbs followed by a period of intermittent overregularization of irregular nouns and verbs. Errors in the model show evidence of phonological conditioning and frequency effects. Furthermore, the network demonstrates a strong tendency to regularize denominal verbs and deverbal nouns and masters the principles of voicing assimilation. Despite being dealt with by a single network, nouns and verbs exhibit some important differences in their profiles of acquisition. Most importantly, noun inflections are acquired earlier than verb inflections. The simulations generate several empirical predictions that can be used to evaluate further the suitability of this type of cognitive architecture in the domain of inflectional morphology, thus pointing the way for close links between computational and empirical research. The model has good secondary data contact and decent input representativeness, but the task veridicality is poor because the task of mapping noun and verb stems to plural and past tense inflections is not likely to play a large role in language acquisition.

Chapter 5, *Finite models of infinite language: A connectionist approach to recursion*, by Morten Christiansen and Nick Chater deals with another theoretical issue that has been seen as strongly supporting a symbolic, rather than a connectionist, approach to language processing: natural language recursion. Since the inception of modern linguistics, there has been considerable emphasis on the importance of recursive phenomena in natural language, and the assumption that any approach to sentence processing must allow for unbounded recursion. Indeed, the existence of different kinds of recursion has had important implications on the choices of symbolic formalisms (e.g., different kinds of generative grammars, different classes of parser and generator) that have been used to explain natural language. From this perspective, natural language recursion presents a difficult challenge to any non-symbolic account of natural language processing.

A range of connectionist approaches have been put forward which attempt to deal with recursion in natural language, although they have not typically achieved the unbounded character of natural language recursion that linguists typically assume. Christiansen and Chater note, though, that the proposition that natural language allows unbounded applications of recursion may make an inappropriate target for connectionist modeling. Instead, they argue that a more appropriate goal for connectionism is to account for the levels of performance that people exhibit when exposed to recursive constructions—to address recursion as a purely psycholinguistic phenomenon, rather than as a linguistic abstraction. Importantly, people's ability to process recursive constructions is quite limited. People produce only a very limited number of complex recursive constructions in naturally occurring speech, and this is reflected in the empirically documented difficulties that people experience when processing such structures.

Christiansen and Chater present a connectionist model of human performance in processing recursive language structures, based on Elman's (1990) Simple Recurrent Network (SRN). The model is trained on simple artificial languages inspired by Chomsky (1957). They find that the qualitative performance profile of the model closely matches human behavior, both on the relative difficulty of center-embedded and cross-dependency, and between the processing of these complex recursive structures and right-branching recursive constructions. Christiansen and Chater analyze how these differences in performance are reflected in the internal representations of the model by performing discriminant analyses on these representations both before and after training. The model has good primary data contact and reasonable task veridicality, but the input representativeness is low because of the abstractness of the artificial languages. More generally, this work suggests a novel explanation of people's limited recursive performance, without assuming the existence of a mentally represented grammar allowing unbounded recursion.

Chapter 6, *Dynamical systems for sentence processing*, by Whitney Tabor and Michael Tanenhaus, like the previous chapter, addresses the question of natural language processing at the level of the sentence, using input patterned on natural language, rather than the more abstract structures used by Christiansen and Chater. Tabor and Tanenhaus suggest that the theory of dynamical systems, originated in the physical sciences, provides a revealing general framework for modeling the representations and mechanisms underlying sentence processing. Recent work in sentence processing (e.g., McRae, Spivey-Knowlton & Tanenhaus, 1998) suggests that parsing biases change fairly continuously over the course of processing the successive words of a sentence. Connectionist networks are good at fitting graded data and their dynamical properties are naturally suited to modeling continuously changing quantities. But the connectionist network that has been most successful in modeling natural language syntax (Elman's SRN, which is used by Christiansen and Chater in the previous chapter) does not explicitly model processing times. They argue that, like many connectionist models at the present time, the SRN is analytically opaque: it is difficult to see the principles underlying its solutions to complex tasks. And it is relativistic—no categorical distinctions are made between grammatical and ungrammatical strings—so it is hard to use linguistic structural insights, which make heavy use of such distinctions, to get past the opaqueness. They suggest that dynamical systems theory, through its insight into the relationship between quantitative and topological properties of systems, offers a solution to these shortcomings.

As in their previous work (Tabor, Juliano, and Tanenhaus, 1997), Tabor and Tanenhaus add a postprocessor to the SRN which has explicit dynamics, thus introducing potentially useful dynamical systems concepts: attractors, basins, saddle-points, trajectories. They call this the Visitation Set Gravitation (VSG) model. Trained on a simple formal language which shares certain key properties with English, the model predicts the important reading time contrasts in a recent study of the real-time evolution of parsing biases (McRae et al., 1998).

Further examination of the VSG model reveals that a standard structural contrast in dynamical systems—between saddle points and attractors—maps onto the fundamental linguistic contrast between ungrammatical and grammatical strings, thus helping to bridge the gap between connectionist models and linguistic theory. And, without further modification of the model, a behaviorally accurate analysis of semantically strange sentences falls out: they are grammatical sentences which involve long trajectories in the state space of the dynamical system and thus have long processing times. This insight helps move work on formal language learning in the much-needed direction of addressing semantic structure.

Overall, the Tabor and Tanenhaus model has good primary data contact and decent task veridicality, but the input representativeness is low because of the simplicity of their formal language. The results suggest that dynamical systems theory is a promising source of ideas for relating the flexible, real-time behavior of the human language processor to its overarching, relatively static, categorical organization. This application of dynamical systems ideas is part of a larger movement within cognitive science (e.g., Kelso, 1997; Port & van Gelder, 1995; Thelen & Smith, 1994), which seeks to understand cognition in dynamical terms. Language processing provides a challenging test case for the application of the dynamical approach, because language has traditionally been conceived from a symbolic perspective. It is interesting, too, to wonder to what extent that connectionist researchers will follow Tabor and Tanenhaus in using ideas from dynamical systems theory to construct and understand connectionist systems. If this does occur, it might represent a substantial departure from the current technical literature on connectionist networks, which is grounded in probability, information theory, and statistical mechanics, rather than dynamical ideas (Bishop, 1995; Frey, 1998).

Chapter 7, *Connectionist models of language production: Lexical access and grammatical encoding* by Gary Dell, Franklin Chang and Zenzi Griffin, moves the focus from how language is understood to how it is produced. Language production has, like language understanding, frequently been characterized as involving the operation of symbolic processes on a generative grammar, and a specification of the message to be produced, in terms of a symbolically encoded underlying "logical form" or "conceptual representation". In contrast to this kind of account, there has also been a long tradition of connectionist theorizing about language production. Indeed, Dell's (1986) "spreading activation" model of speech production was one of the most important models in the revival of interest in connectionist models of psychological processes, which began in the early to mid-nineteen-eighties. In their chapter, Dell et al. describe the most recent developments in this approach to modeling speech production. Specifically, they outline the issues and challenges that must be addressed by connectionist models of lexical access and grammatical encoding, and review three recent models. The models illustrate the value of a spreading activation approach to lexical access in production, the need for sequential output in both phonological and grammatical encoding, and the potential for accounting for structural effects on phonological errors and syntactic priming. These models account for a broad range of data on speech production, from the analysis of speech errors, to the performance of aphasic patients, and to results from syntactic priming studies. Indeed, in speech production research, the interplay between connectionist modeling and the gathering of empirical data that we view as constitutive of connectionist psycholinguistics is particularly well-developed.

Dell et al. consider several specific models, rather than attempting a single overarching model of speech production. Individually, the models have good primary or secondary data contact and good task veridicality, but all models suffer from relatively low input representativeness because the models only cover small language fragments. An important question for future research concerns the degree to which models of specific aspects of speech production can be integrated in a cohesive way, an issue that also arises in relation to the models of speech and language processing described in earlier chapters of this book.

The chapters so far have focused on the comprehension and production of speech, rather than how written language is processed. The reading of single words has, in particular, been an area of intense connectionist research. Chapter 8, *A connectionist approach to word reading and acquired dyslexia: Extension to sequential processing*, by David Plaut, outlines a new model of reading, building on the long research tradition. Plaut begins by discussing some general principles of the connectionist approach to word reading—of which he is a leading proponent—including distributed representation, graded learning of statistical structure, and interactivity in processing. These principles have led to the development of explicit computational models that account for an impressively broad range of data, from the interaction of word frequency and spelling-sound consistency in normal skilled reading, to analogous effects in the reading errors of surface dyslexic patients, and the co-occurrence of visual and semantic errors in deep dyslexia.

Plaut notes, though, that there have been recent empirical challenges to these models, and the approach in general, relating to the influence of orthographic length on the naming latencies of both normal and dyslexic readers. For instance, the models account for relatively little variance associated with individual words in large databases of naming latencies, partly due to insufficient sensitivity to orthographic length. The models also underestimate length effects in the naming latencies for nonwords. This kind of empirical challenge is an illustration of the productive interaction between connectionist modeling and empirical research—predictions of connectionist models have had a crucial impact in directing the search for relevant empirical confirmation or disconfirmation.

Plaut addresses this challenge by presenting a new model that generates sequential phonological output in response to written input. He trains an SRN (Elman, 1990) to produce sequences of single phonemes as output when given position-specific letters as input. The model was also trained to maintain a representation of its current position within the input string. When the model found a peripheral portion of the input difficult to pronounce, it used the position signal to refixate the input, shifting the peripheral portion to the point of fixation where the model has had more experience in generating pronunciations. In this way, the model could apply the knowledge tied to the units at the point of fixation to any difficult portion of the input. Early on in training, the model required multiple fixations to read words, but as the model became more competent it eventually read most words in a single fixation. The model could also read nonwords about as well as skilled readers, occasionally falling back on a refixation strategy for difficult nonwords. The model exhibits an effect of orthographic length and a frequency-by-consistency interaction in its naming latencies. When subject to peripheral damage, the model exhibits an increased length effect which interacts with word frequency, characteristic of letter-by-letter reading in pure alexia. The model thus has a good primary data contact and good task veridicality, but input representativeness suffers somewhat because the model is only trained on monosyllabic words. Plaut notes that the model is not intended as a fully adequate account of all the relevant empirical phenomena. But the model provides a compelling demonstration of how connectionist models may be extended to provide deeper insight into sequential processes in reading.

Plaut's chapter concludes the first part of the book, which reviews current models of connectionist language processing. The second part of the book consists of three insightful perspectives on the significance, interpretation and utility of connectionist psycholinguistics, by eminent researchers in the cognitive science of language processing.

Chapter 9, *Constraint satisfaction in language acquisition and processing* by Mark Seidenberg and Maryellen MacDonald sets out the most radical connectionist agenda, seeing connectionism as potentially undermining classical symbolic theorizing in linguistics and psycholinguistics. In particular, they see connectionist psycholinguistics as part of a larger theoretical framework focusing on probabilistic constraints on language processing and language acquisition (Seidenberg, 1997). This probabilistic framework offers an alternative viewpoint on language and language use to that found in generative linguistics. The generative approach attempts to characterize knowledge of language (i.e., competence grammar) and then asks how this knowledge is acquired and used. Seidenberg and MacDonald's probabilistic approach is performance oriented: the goal is to explain how people comprehend and produce utterances and how children acquire this skill. From a probabilistic perspective, using language is thought to involve exploiting multiple probabilistic constraints over various types of linguistic and nonlinguistic information.

Children begin accumulating this information at a young age. The same constraint-satisfaction processes that are central to language use in adulthood also serve as the bootstrapping processes that provide entry into language in childhood. Framing questions about acquisition in terms of models of skilled language use has important consequences for arguments concerning language learnability and holds out the possibility of a unified theory of acquisition and use.

Seidenberg and MacDonald put forward a vigorous case for opposition between connectionist and symbolic approaches to language. But this is, of course, by no means the only possible viewpoint. Language might, instead, be viewed as emerging from a mixture of linguistic rules, which can be specified in symbolic terms, and probabilistic factors that determine how language is used in specific contexts; and, indeed, symbolic linguistic rules need not, perhaps, be quite as rigid as is typically assumed. Thus, a more conciliatory line between connectionist psycholinguistics and symbolic, generative linguistics may be imagined.

Chapter 10, *Grammar-based connectionist approaches to language*, by Paul Smolensky, outlines a specific conception of how connectionist and symbolic theorizing about language might be integrated, rather than set against each other. In particular, Smolensky argues that connectionist research on language can and must involve the development of grammar formalisms rather than merely producing connectionist computer models. From formulations of the fundamental theoretical commitments of connectionism and of generative grammar, it is argued that these two paradigms are mutually compatible: the commitments of connectionism concern computational principles, and those of generative grammar concern explanations of certain fundamental empirical characterizations of human language. Integrating the basic assumptions of the two paradigms results in formal theories of grammar that centrally incorporate a certain degree of connectionist computation. Two such grammar formalisms—Harmonic

Grammar (Legendre, Miyata and Smolensky, 1990) and Optimality Theory (Prince and Smolensky, 1997)—are briefly introduced to illustrate grammar-based approaches to connectionist language research. The strengths and weaknesses of grammar-based research and more traditional model-based research are argued to be complementary: grammar-based research more readily allows explanations of general patterns of language, while model-based research more readily allows exploration of the full richness and power of connectionist computational mechanisms. This complementarity of strengths suggests a significant role for both strategies in the spectrum of connectionist language research.

Smolensky's standpoint provides a counterweight to the view that connectionist psycholinguistics should attempt to overturn previous theorizing about language and language processing. Moreover, the synthesis of principles from connectionism and generative grammar that he outlines gains considerable credibility from its very widespread influence in modern linguistics. Indeed, optimality theory is widely viewed within linguistics as one of the central theoretical developments within the field. Smolensky's approach may not, however, satisfy more radical connectionists, who may see the move from specific and implementable connectionist models of psychological processes (such as are described in the first half of this book) to abstract connectionist principles as too great a departure from the original aims of the connectionist paradigm.

Chapter 11, *Connectionist sentence processing in perspective*, by Mark Steedman presents an outsider's perspective on the project of connectionist psycholinguistics. Steedman has been associated with symbolic approaches to language, and has been involved in pioneering novel linguistic formalisms, such as categorial grammar, as well as carrying out highly influential computational and experimental work on human language processing. Steedman focuses on connectionist sentence processing, a topic discussed in Chapters 5 and 6 of this book. Steedman argues that the emphasis in the connectionist sentence-processing literature on distributed representation and emergence of grammar from such systems seems to have prevented connectionists and symbolic theorists alike from recognizing the often close relations between their respective systems. He argues that SRN models (Elman, 1990) are more directly related to stochastic Part-of-Speech taggers than to parsers or grammars as such, while recursive auto-associative memory of the kind pioneered by Pollack (1990) and incorporated in many hybrid connectionist parsers since may be useful for grammar induction from a network-based conceptual structure as well as for structure-building.

These observations suggest some interesting new directions for connectionist sentence processing research, including more efficient devices for representing finite state machines, and acquisition devices based on a distinctively connectionist-grounded conceptual structure. Thus, Steedman, like Smolensky, argues for an integration of connectionist and symbolic views of language and language processing. But Steedman and Smolensky differ concerning the nature of the integration. Whereas Smolensky argues that connectionist principles should be integrated into grammar formalisms, Steedman sees connectionist networks as integrating with symbolic language processing mechanisms, to produce a hybrid computational account of language processing and acquisition. And both Smolensky and Steedman differ from the more radical agenda of Seidenberg and MacDonald, which aims to replace, rather than interconnect with, previous theories of language processing and structure. Clearly, only the future development of connectionist research will decide which of these perspectives, each of which is persuasively argued, proves to be the most fruitful.

The Significance of Connectionist Psycholinguistics

Current connectionist models involve important simplifications with respect to real natural language processing. In some cases, these simplifications are relatively modest. For example, models of reading aloud typically ignore how eye movements are planned and how information is integrated across eye-movements; they also tend to ignore the sequential character of speech output, and typically deal only with short words. In other cases, the simplifications are more drastic. For example, connectionist models of syntactic processing involve vocabularies and grammars that are vastly simplified. However, it is important to note that symbolic models in many cases have lower task veridicality and input representativeness than their connectionist counterparts. Furthermore, many symbolic models may give the appearance of good data contact simply because they have not yet been implemented and have therefore not been tested in an empirically rigorous way, in contrast to the connectionist models.

The present breadth and significance of connectionist psycholinguistics, as evidenced by the chapters in this volume, suggests that the approach has considerable potential. Despite some attempts to argue for a priori limitations on connectionist language processing (e.g., Pinker & Prince, 1988; Marcus, 1998), connectionist psycholinguistics has already had a major impact on the psychology of language.

First, connectionist models have provided the first fully explicit and psychologically relevant computational models in a number of language processing domains, such as reading and past tense learning. Previous accounts in these areas consisted of "box-and-arrow" flow diagrams rather than detailed computational mechanisms. Whatever the lasting value of connectionist models themselves, they have certainly raised the level of theoretical debate in these areas, by challenging theorists of all viewpoints to provide computationally explicit accounts.

Second, the centrality of learning in connectionist models has brought a renewed interest in mechanisms of language learning (Bates & Elman, 1993). While Chomsky (e.g., 1986) has argued that there are "universal" aspects of language which are innate, the vast amount of information specific to the language that the child acquires must be learned. Connectionist models provide mechanisms for how (at least some of) this learning might occur, whereas previous symbolic accounts of language processing have not taken account of how learning might occur. Furthermore, the attempt to use connectionist models to learn syntactic structure encroaches on the area of language that Chomsky has argued innate information must be central. The successes and failures of this program thus directly bear on the validity of this viewpoint.

Third, the dependence of connectionist models on statistical properties of their input has been one contributing factor in the upsurge of interest in the role of statistical factors in language learning and processing (MacWhinney, Leinbach, Taraban & McDonald, 1989; Redington & Chater, 1998). This renewed interest in the statistical properties of language and statistical methods of analysis is, of course, entirely compatible with the view that language processing takes account of structural properties of language, as described by classical linguistics. But more radical connectionists have, as we have noted, also attempted to encroach on the territory of classical linguistics.

Finally, connectionist systems have given rise to renewed theoretical debate concerning what it really means for a computational mechanism to implement a rule, whether there is a distinction between "implicit" and "explicit" rules (see e.g., Davies, 1995, for discussions), and which kind should be ascribed to the human language processing system.

The potential implications of a realistic connectionist approach to language processing are enormous. If realistic connectionist models of language processing can be provided, then the possibility of a radical rethinking not just of the nature of language processing, but of the structure of language itself, may be required. It might be that the ultimate description of language resides in the structure of complex networks, and can only be approximately expressed in terms of structural rules, in the style of generative grammar (Seidenberg & MacDonald, Chapter 9). On the other hand, it may be that connectionist models can only succeed to the extent that they build in standard linguistic constructs (Smolensky, Chapter 10), or fuse with symbolic models to create a hybrid approach (Steedman, Chapter 11). We suggest that the only way to determine the ultimate value of connectionist psycholinguistics is to pursue it with the greatest possible creativity and vigor, as exemplified by the chapters in this volume.

References

Bates, E.A. & Elman, J.L. (1993). Connectionism and the study of change. In M.J. Johnson (Ed.), *Brain development and cognition* (pp. 623–642). Cambridge, MA: Basic Blackwell.

Berent, I., Pinker, S. & Shimron, J. (1999). Default nominal inflection in Hebrew: Evidence for mental variables. *Cognition*, *72*, 1–44.

Bishop, C.M. (1995). *Neural networks for pattern recognition.* New York: Oxford University Press.

Chomsky, N. (1957). *Syntactic structures.* The Hague: Mouton.

Chomsky, N. (1986). *Knowledge of language.* New York: Praeger.

Christiansen, M.H. & Chater, N. (2000). *Connectionist Psycholinguistics: Capturing the empirical data.* Manuscript submitted for publication.

Coltheart, M., Curtis, B., Atkins, P. & Haller, M. (1993). Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review*, *100*, 589–608.

Davies, M. (1995). Two notions of implicit rules. In J.E. Tomberlin (Ed.), *Philosophical Perspectives, Volume 9: AI, Connectionism, and Philosophical Psychology.* Atascadero, CA: Ridgeview Publishing Company.

Dell, G.S. (1986). A spreading activation theory of retrieval in language production. *Psychological Review*, *93*, 283–321.

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

Frey, B.J. (1998). *Graphical models for machine learning and digital communication.* Cambridge, MA: MIT Press.

Gaskell, M.G. & Marslen-Wilson, W.D. (1999). *Representation and competition in the processing of spoken words.* Manuscript submitted for publication.

Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, *68*, 1–76.

Hoeffner, J. (1997). *Are rules a thing of the past? A single mechanism account of English past tense acquisition and processing.* Unpublished Ph.D. dissertation. Pittsburgh, PA: Carnegie Mellon University.

Just, M.A. & Carpenter, P.A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, *98*, 122–149.

Kelso, J.A.S. (1997). *Dynamic patterns: The self-organization of brain and behavior.* Cambridge, MA: MIT Press.

Legendre, G., Miyata, Y. & Smolensky, P. (1990). Harmonic grammar – a formal multi-level connectionist theory of linguistic well-formedness: Theoretical foundations. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 388–395). Hillsdale, NJ: Lawrence Erlbaum.

MacWhinney, B., Leinbach, J., Taraban, R. & McDonald, J. (1989). Language learning: Cues or rules? *Journal of Memory and Language*, *28*, 255–277.

Marcus, G.F. (1998). Rethinking eliminative connectionism. *Cognitive Psychology*, *37*, 243–282.

McRae, K., Spivey-Knowlton, M.J. & Tanenhaus, M.K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, *38*, 282–312.

Pinker, S. (1991). Rules of language. *Science*, *253*, 530–535.

Pinker, S. & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, *28*, 73–193.

Pollack, J.B. (1990). Recursive distributed representations. *Artificial Intelligence*, *46*, 77–105.

Port, R.F. & van Gelder, T. (Eds.) (1995). *Mind as motion: Explorations in the dynamics of cognition.* Cambridge, MA: MIT Press.

Prince, A. & Smolensky, P. (1997). Optimality: From neural networks to universal grammar. *Science*, *275*, 1604–1610.

Redington, M. & Chater, N. (1998). Connectionist and statistical approaches to language acquisition: A distributional perspective. *Language and Cognitive Processes*, *13*, 129–191.

Rumelhart, D.E. & McClelland, J.L. (1986). On learning of past tenses of English verbs. In McClelland, J.L. & Rumelhart, D.E. (Eds.), *Parallel distributed processing, Volume 2* (pp. 216–271). Cambridge, MA: MIT Press.

Seidenberg, M.S. (1997). Language acquisition and use: Learning and applying probabilistic constraints. *Science*, *275*, 1599–1603.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, *11*, 1–74.

Tabor, W., Juliano, C. & Tanenhaus, M.K. (1997). Parsing in a dynamical system: An attractor-based account of the interaction of lexical and structural constraints in sentence processing. *Language and Cognitive Processes*, *12*, 211–271.

Thelen, E. & Smith, L.B. (1994). *A dynamic systems approach to the development of cognition and action.* Cambridge, MA: MIT Press.

Author Note

We are grateful to the Cognitive Science Society for allowing us to put this this book together based on papers that appeared in *Cognitive Science*. All of the chapters except one are revised and updated versions of earlier solicited papers found in the special issue on “Connectionist models of human language processing: Progress and prospects”, Vol. 23(4), 1999, edited by Morten Christiansen, Nick Chater and Mark Seidenberg. Chapter 5 is a substantially revised version of an article that appeared in Vol. 23(2), 1999. We would like to thank the contributors for updating and revising their papers for this book, and to the reviewers for helping to ensure a very high quality of papers throughout.

This work was partially supported by the Leverhulme Trust and by European Commission grant RTN - HPRN-CT-1999-00065 to NC.

Footnotes