Connectionist Approaches to Natural Language Processing

Ronan G. Reilly and Noel E. Sharkey, editors
(University College Dublin and University of Exeter)

Reviewed by
James Henderson
University of Pennsylvania

With the growing popularity of statistical approaches to natural language processing, it is natural to consider how connectionist techniques for learning and using soft constraints might be applicable to this area. However, nagging questions about the representational and computational adequacy of connectionist networks for processing language make the fruitfulness of this endeavor uncertain. This book addresses both these issues. It contains background and overview sections by the editors, plus 13 chapters covering a broad spectrum of recent work involving connectionist models and natural language processing. Each chapter can be read on its own, or the book can be read in its entirety. A casual familiarity with connectionist networks is adequate for understanding most of the book.

The book begins with an introductory chapter that includes some background about connectionist natural language processing and an overview of the other chapters. Many references are given to earlier work, and the discussion provides a good context for the more recent work presented in the book. The remaining chapters are organized into four sections: semantics, syntax, representational adequacy, and computational psycholinguistics. Each section has its own introduction, plus chapters written by researchers in the area.

The section on connectionist approaches to semantics includes four chapters that range from lexical semantic representations to "story comprehension." In Chapter 2, Dyer, Flowers, and Wang propose a connectionist architecture for learning representations for symbols by training on the relationships between the symbols. They apply this method to encoding semantic networks using one symbol per semantic node. In Chapter 3, Sutcliffe argues for representing word meanings using an unstructured set of semantically interpretable features with associated "centrality" values. This representation is uniform across words and provides a simple measure of the distance between two concepts. Wermter and Lehnert address the issue of disambiguating PP attachment in noun phrases using semantic information. They train one set of networks to produce plausibility values for individual attachments and use a constraint satisfaction network to combine these preferences with structural constraints to determine the best set of attachments. St. John and McClelland apply connectionist learning techniques and a constraint satisfaction network to a variety of tasks involving semantic preferences. They demonstrate the usefulness of these techniques for some problems, but their claim that no additional mechanisms are needed for sentence and story comprehension is not substantiated.

The section on syntax contains two chapters that use hand-structured networks to implement symbolic grammars of the familiar kind. In Chapter 6, Rager addresses parsing and correcting extragrammatical input. Individual constraints are represented in constraint satisfaction (winner-take-all) subnetworks, and in the event of an in-

consistency the best alternative is chosen. Schnelle and Doust implement an Earley parser for context-free grammars. This article was hard to follow, but it seemed to simply illustrate that just about anything can be compiled into units and links given a sufficiently broad definition of those primitives.

The section on representational adequacy presents three perspectives on the role of symbolic representations in connectionist models. Bever argues that connectionist models can only capture habitual, associative knowledge, but that this ability makes them an important tool for use within explicit structural symbolic theories. He argues that connectionist models that claim to circumvent the need for predefined knowledge actually have such knowledge built in, and describes a network built with explicitly predefined structured knowledge that learns to segment utterances into linguistically appropriate phrases. Berg discusses two ways in which connectionist models of the time were inadequate, namely distinguishing between multiple instances of the same concept and creating new concepts quickly. He proposes a hybrid connectionist/marker-passing model to address these problems while staying massively parallel. Dorffner argues that a model of language need not and should not depend on explicit representations specified by a researcher. This harkens back to the debate over whether linguistic behavior can be explained through a general-purpose learning mechanism that has no prespecified (innate) linguistic knowledge (cf Chomsky 1966).

The last section is on computational psycholinguistics. The first two chapters propose alternatives to McClelland and Elman's (1986) TRACE model of speech perception, where feedback from the word level facilitates recognition at the phoneme level. Massaro proposes a feedforward (bottom-up) model of speech perception phenomena based on fuzzy logic. The argument would be more convincing if the network weren't trained on the data to be explained. Norris argues for a feedforward network with a single hidden layer that is shared by two unconnected output layers for word recognition and phoneme recognition. He shows that the apparent top-down effects can be the result of the top-down aspects of the backpropagation learning algorithm, and do not have to result from top-down information during processing. O'Seaghdha, Dell, Peterson, and Juliano discuss two connectionist models of form-related priming in lexical access, and opt for the one with articulated representations of the prime and target words. Mozer and Behrmann model a reading disorder caused by brain damage as an impairment in the attention component of a previously proposed connectionist model of word recognition in reading.

Research on connectionist natural language processing is still in its early stages, and this is evident in this book. Consequently, readers of *Computational Linguistics* may be disappointed with the preliminary nature of many of the models compared with more traditional approaches to language. This also accounts for this book's emphasis on connectionist approaches rather than natural language processing. Although most of the questions and proposals discussed in this book will be with us for a long time to come, it is inevitable that a book will not include the most recent developments in such a quickly changing field. For natural language processing, the most important of these is the development of architectures that support symbolic processing without violating the basic tenets of connectionism. For a comprehensive bibliography on this topic, see Sun and Bookman (1993). Despite these inherent problems for any book on connectionist natural language processing, this book is worthwhile reading for anyone interested in the lively debate over the adequacy of connectionist models, in the old debate over learning versus prespecified knowledge, in the application of connectionist models to problems involving soft constraints, or in connectionist models in general.

References

Chomsky, Noam (1966). *Cartesian Linguistics*. Harper and Row.

McClelland, J. L., and Elman, J. L. (1986). "The TRACE model of speech perception." *Cognitive Psychology*, 18, 1–86.

Sun, R., and Bookman, L. (editors) (1993). *Computational Architectures Integrating Neural and Symbolic Processes*. Kluwer Academic Publishers.

James Henderson is a Ph.D. candidate in Computer and Information Science at the University of Pennsylvania, majoring in Computational Linguistics. His dissertation is on syntactic parsing in a recently proposed connectionist architecture that supports symbolic computation. Henderson's address is: Institute for Research in Cognitive Science, 3401 Walnut St, Philadelphia, PA 19104; email: henders@linc.cis.upenn.edu.