SARDSRN: A Neural Network Shift-Reduce Parser

Marshall R. Mayberry, III and Risto Miikkulainen
Department of Computer Sciences
The University of Texas
Austin, TX 78712, U.S.A
(martym,risto@cs.utexas.edu)

Abstract

Simple Recurrent Networks (SRNs) have been widely used in natural language tasks. SARDSRN extends the SRN by explicitly representing the input sequence in a SARDNET self-organizing map. The distributed SRN component leads to good generalization and robust cognitive properties, whereas the SARDNET map provides exact representations of the sentence constituents. This combination allows SARDSRN to learn to parse sentences with more complicated structure than can the SRN alone, and suggests that the approach could scale up to realistic natural language.

1 Introduction

The subsymbolic approach (i.e. neural networks with distributed representations) to processing language is attractive for several reasons. First, it is inherently robust: the distributed representations display graceful degradation of performance in the presence of noise, damage, and incomplete or conflicting input [Miikkulainen, 1993; St John and McClelland, 1990]. Second, because computation in these networks is constraint-based, the subsymbolic approach naturally combines syntactic, semantic, and thematic constraints on the interpretation of linguistic data [McClelland and Kawamoto, 1986]. Third, subsymbolic systems can be lesioned in various ways and the resulting behavior is often strikingly similar to human impairments [Miikkulainen, 1993;. 1996; Plant, 19911 These properties of subsymbolic systems have attracted many researchers in the hope of accounting for interesting cognitive phenomena, such as role-binding and lexical errors resulting from memory interference and overloading, aphasic and dyslexic impairments resulting from physical damage, and biases, defaults and expectations emerging from training history [Miikkulainen, 1997; 1996; 1993; Plaut and Shallice, 1992].

Since its introduction in 1990, the simple recurrent network (SRN) [Elman, 1990] has become a mainstay in connectionist natural language processing tasks such as lexical disambiguation, prepositional phrase attachment, active-passive transformation, anaphora resolution, and translation [Allen, 1987; Chalmers, 1990; Munro *et al.,*1991; Touretzky, 1991]. This paper describes an extension to the standard SRN, which utilizes SARDNET [James and Miikkulainen, 1995], a self-organizing map algorithm designed to represent sequences. SARDNET permits the sequence information to remain explicit, yet distributed in the sense that similar sequences result in similar patterns on the map. SARDSRN, THe combination of the SRN and SARDNET, effectively solves the fundamental memory accuracy limitations of the SRN, and allows the processing of sentences of realistic length.

This paper shows how SARDSRN improves upon the performance of the SRN in a nontrivial syntactic shift-reduce parsing task. The results show that SARDSRN outperforms the SRN in this task by providing an effective solution to the memory problem. SARDSRN therefore forms a solid foundation for building a subsymbolic parser of realistic language.

2 The Task: Shift-Reduce Parsing

The task taken up in this study, shift-reduce (SR) parsing, is one of the simplest approaches to sentence processing that nevertheless has the potential to handle a substantial subset of English [Tomita, 1986]. Its basic formulation is based on the pushdown automata for parsing context-free grammars, but it can be extended to context-sensitive grammars as well.

The parser consists of two data structures: the input buffer stores die sequence of words remaining to be read, and the partial parse results are kept on the stack (figure 1). Initially the stack is empty and the entire sentence is in the input buffer. At each step, the parser has to decide whether to shift a word from the buffer to the stack, or to reduce one or more of the top elements of the stack into a new element representing their combination. For example, if the top two elements are currendy *NP* and VP_t the parser reduces diem into 5, corresponding to the grammar rule $S \rightarrow NP\ VP$ (step 17 in figure 1). The process stops when the elements in the stack have been reduced to S, and no more words remain in the input "The reduce actions performed by the parser in this process constitute the parse result, such as the syntactic parse tree (line 18 in figure 1).

The sequential scanning process and incremental forming of partial representations is a plausible cognitive model for language understanding. SR parsing is also very efficient, and lends itself to many extensions. For example, the parse rules can be made more context sensitive by taking more of the stack and the input buffer into account Also, the partial parse results may consist of syntactic or semantic structures.

Stack	Input Buffer	Action	
	the boy who liked the girl chased the cat .	1	Shift
the	boy who liked the girl chased the cat .	2	Shift
the boy	who liked the girl chased the cat .	3	Reduce
NP[the,boy]	who liked the girl chased the cat .	4	Shift
NP[the,boy] who	liked the girl chased the cat .	5	Shift
NP[the,boy] who liked	the girl chased the cat .	6	Shift
NP[the,boy] who liked the	girl chased the cat .	7	Shift
NP[the,boy] who liked the girl	chased the cat .	8	Reduce
NP[the,boy] who liked NP[the,girl]	chased the cat .	9	Reduce
NP[the,boy] who VP[liked,NP[the,girl]]	chased the cat .	10	Reduce
NP[the,boy] RC[who,VP[liked,NP[the,girl]]]	chased the cat .	11	Reduce
NP[NP[the,boy],RC[who,VP[liked,NP[the,girl]]]]	chased the cat .	12	Shift
NP[NP[the,boy],RC[who,VP[liked,NP[the,girl]]]] chased	the cat .	13	Shift
NP[NP[the,boy],RC[who,VP[liked,NP[the,girl]]]] chased the	cat .	14	Shift
NP[NP[the,boy],RC[who,VP[liked,NP[the,girl]]]] chased the cat	.	15	Reduce
NP[NP[the,boy],RC[who,VP[liked,NP[the,girl]]]] chased NP[the,cat]	.	16	Reduce
NP[NP[the,boy],RC[who,VP[liked,NP[the,girl]]]] VP[chased,NP[the,cat]]	.	17	Reduce
S[NP[NP[the,boy],RC[who,VP[liked,NP[the,girl]]]],VP[chased,NP[the,cat]]]		18	Stop

Figure 1: Shift-Reduce Parang a Sentence. Each step in the parse is represented by a line fro The current stack is at left, the input buffer in the middle, and the parsing decision in the current situation at right. At each step, the parser either shifts a word onto the stack, or reduces the top elements of the stack into a higher-level representation, such as the $\rightarrow$ NP[the,boy] (step 3). (Phrase labels such as "NP" and "RC" are only used in this figure to make the process clear.)

The general SR model can be implemented in many ways. A set of symbolic shift-reduce rules can be written by hand or learned from input examples [Hermjacob and Mooney, 1997; Simmons and Yu, 1991; Zelle and Mooney, 1996]. It is also possible to train a neural network to make parsing decisions based on the current stack and the input buffer. If trained properly, the neural network can generalize well to new sentences [Simmons and Yu, 1992]. Whatever correlations there exist between the word representations and the appropriate shiftz/reduce decisions, the network will learn to utilize them.

Another important extension is to implement the stack as a neural network. This way the parser can have access to the entire stack at once, and interesting cognitive phenomena in processing complex sentences can be modeled. The SPEC system [Miikkulainen, 1996] was a first step in this direction. The stack was represented as a compressed distributed representation, formed by a RAAM (Recursive Auto-Associative Memory) auto-encoding network [Pollack, 1990]. The resulting system was able to parse complex relative clause structures. When the stack representation was artificially lesioned by adding noise, the parser exhibited very plausible cognitive performance. Shallow center embeddings were easier to process, as were sentences with strong semantic constraints in the role bindings. When the parser made errors, it usually switched the roles of two words in the sentence, which is what people also do in similar situations. A symbolic representation of the stack would make modeling such behavior very difficult.

The SPEC architecture, however, was not a complete implementation of SR parsing; it was designed specifically for embedded relative clauses. For general parsing, the stack needs to be encoded with neural networks to make it possible to parse more varied linguistic structures. We believe that the generalization and robustness of subsymbolic neural networks will result in powerful, cognitively valid performance. However, the main problem of limited memory accuracy of the SEN parsing network must first be solved. An architecture that will do that, SARDSRN, will be described next

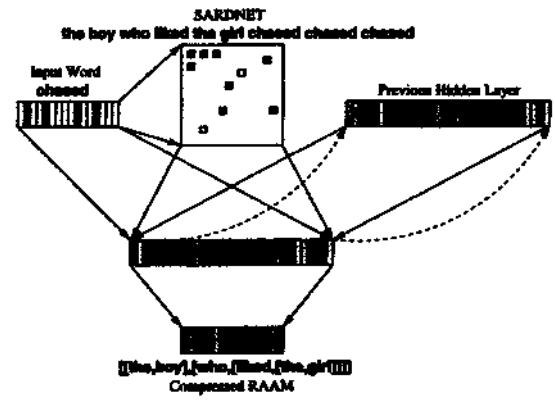

Figure 2: The SARDSRN Network. This snapshot shows the network during step 11 of figure 1. The representation for the current input word, chased, is shown at top left. Each word is input to the SARDNET map, which builds a representation for the sequence word by word. At each step, the previous activation of the hidden layer is copied (as indicated by the dotted line) to the Previous Hidden Layer assembly. This activation, together with the current input word and the current SARDNET pattern, is propagated to the hidden layer of the SRN network. As output, the network generates the compressed RAAM representation of the top element in the shift-reduce stack at this state of the parse (in this case, line 12 in figure 1). SARDNET is a map of word representations, and is trained through the Self-Organizing Map (SOM) algorithm [Kohonen, 1995; 1990]. All other connections are trained through backpropagation.

3 The SARDSRN parser architecture

3.1 Simple Recurrent Network

The starting point for SARDSRN (figure 2) is the simple recurrent network. The network reads a sequence of input word representations into output patterns representing the parse results, such as syntactic or case-role assignments for the words. At each time step, a copy of the hidden layer is saved and used as input during the next step, together with the next word. In this way each new word is interpreted in the context of the entire sequence so far, and die parse result is gradually formed at the output.

The SRN architecture can be used to implement a shift-reduce parser in the following way: the network is trained to step through the parse (such as that in figure l), generating a compressed distributed representation of the top element of the stack at each step (formed by a RAAM network: section 4.1). The network reads the sequence of words one word at a time, and each time either shifts the word onto the stack (by passing it through the network, e.g. step 1), or performs one or more reduce operations (by generating a sequence of compressed representations corresponding to the top element of the stack: e.g. steps 8-11). After the whole sequence is input, the final stack representation is decoded into a parse result such as a parse tree. Such an architecture is powerful for two reasons: (1) During the parse, the network does not have to guess what is coming up later in the sentence, as it would if it always had to shoot for the final parse result; its only task is to build a representation of the current stack in its hidden layer and the top element in its output. (2) Instead of having to generate a large number of different stack states at the output, it only needs to output representations for a relatively small number of common substructures. Both of these features make learning and generalization easier.

A well-known problem with the SRN model is its low memory accuracy. It is difficult for it to remember items that occurred several steps earlier in the input sequence, especially if the network is not required to produce mem in the output layer during the intervening steps [Stolcke, 1990; Miikkulainen, 1996]. The intervening items are superimposed in the hidden layer, obscuring the traces of earlier items. Nor has simply increasing the size of the hidden layer or lowering the learning rate been found to offer much advantage. As a result, parsing with an SRN has been limited to relatively simple sentences with shallow structure.

3.2 SARDNET

The solution described in this paper is to use an explicit representation of the input sequence as additional input to the hidden layer. This representation provides more accurate information about the sequence, such as the relative ordering of the incoming words, and it can be combined with the hidden layer representation to generate accurate output that retains all the advantages of distributed representations. The sequence representation must be explicit enough to allow such cleanup, but it must also be compact and generalize well to new sequences.

The SARDNET (Sequential Activation Retention and Decay Network) [James and Miikkulainen, 1995] self-organizing map for sequences has exactly these properties. SARDNET is based cm the Self-Organizing Map neural network [Kohonen, 1990; 1995], and organized to represent die space of all possible word representations. As in a conventional self-organizing map network, each input word is *mapped* onto a particular map node called the maximally-responding unit, or winner. The weights of the winning unit and all the nodes in its neighborhood are updated according to the standard adaptation rule to better approximate the current input The size of the neighborhood is set at the beginning of the training and reduced as the map becomes more organized.

In SARDNET, the sentence is represented as a distributed activation pattern on the map (figure 2). For each word, the maximally responding unit is activated to a maximum value of 1.0, and the activations of units representing previous words are decayed according to a specified decay rate (e.g. 0.9). Once a unit is activated, it is removed from competition and cannot represent later words in the sequence. Each unit may then represent different words depending on the context, which allows for an efficient representation of sequences, and also generalizes well to new sequences.

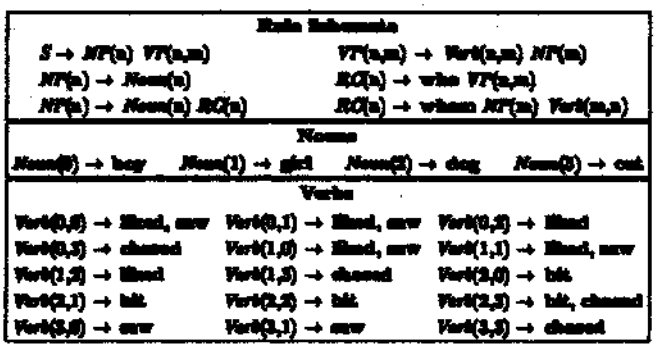

Figure 3: Grammar. This phrase structure grammar generates sentences with subject- and object-extracted relative clauses. The rule schemata with noun and verb restrictions ensure agreement between subject and object depending on the verb in the clause. Lexicon items are given in bold face,

In the SARDSRN architecture, SARDNET is used to directly handle the memory limitation of the SRN. A SARDNET representation of the input sentence is formed at the same time as the SRN hidden layer representation, and used together with the previous hidden layer representation and the next word as input to the hidden layer (figure 2). This architecture allows the SRN to perform its task with significantly less memory degradation. The sequence information remains accessible in SARDNET, and the SRN is able to focus on capturing correlations relating to sentence constituent structure during parsing.

4 Experiments

4.1 Input Data, Training, and System Parameters

The data used to train and test the SRN and SARDSRN networks were generated from the phrase structure grammar in figure 3, adapted from a grammar that has become common in the literature [Elman, 1991; Miikkulainen, 1996]. Since our focus was on shift-reduce parsing, and not processing relative clauses per se, sentence structure was limited to one relative clause per sentence. From this grammar training targets corresponding to each step in the parsing process were obtained. For shifts, the target is simply the current input. In these cases, the network is trained to auto-associate, which these networks are good at. For reductions, however, the targets consist of representations of the partial parse trees that result from applying a grammatical rule. For example, the reduction of the sentence fragment who liked the girl would produce the partial parse result [who,[liked,[the^iri]]]. Two issues arise: how should the parse trees be represented, and how should reductions be processed during sentence parsing?

the whom	10000000 01100000	who	01010000 11111111
boy girl	00101000 00100100	dog cat	00100010 00100001
chased liked	00011000 00010100	saw 1 bit	00010010 00010001

Figure 4: Lexicon, Each word representation is put together from a part~of-speech identifier (first four components) and a unique ID tag Oast four). This encoding is then repeated eight times to form a 64-unit word representation. Such redundancy makes it easier to identify the word.

The approach taken in this paper is the same as in SPEC (section 2), as well as in other connectionist parsing systems [Miikkulainen, 1996; Berg, 1992; Sharkey and Sharkey, 1992]. Compressed representations of all the partial syntactic parse trees using RAAM are built up through auto-association of die constituents. This training is performed beforehand separately from the parsing task. Once formed, the compressed representations can be decoded into their constituents using just the decoder portion of the R AAM architecture.

In shift-reduce parsing, the input buffer after each "Reduce" action is unchanged; rather, the reduction occurs on the stack. Therefore, if we want to perform the reductions one step at a time, the current word must be maintained in the input buffer until the next "Shift" action. Accordingly, the input to the network consists of the sequence of words that make up the sentence with the input word repeated for each reduce action, and the target consists of representations of the top element of the stack (as shown in figure 1).

Word representations were hand-coded to provide basic part-of-speech information together with a unique ID tag that identified the word within the syntactic category (figure 4). The basic encoding of eight units was repeated eight times to fill out a 64-unit representation, The 64-unit representation length was needed to encode all of the partial parse results formed by RAAM, and redundancy in the lexical items facilitate learning.

Four data sets of 20%, 40%, 60%, and 80% of the 436 sentences generated by the grammar were randomly selected to train both parsers, and each parser was trained on each dataset four times. Training on all thirty-two runs was stopped when die error on a 22-sentence (5%) validation set began to level off. The same validation set was used for all the simulations and was randomly drawn from a pool of sentences that did not appear in any of the training sets. Testing was then performed on the remaining sentences that were neither in the training set nor in the validation set.

The SRN network architecture consisted of a 64-unit input layer, 200-unit hidden and context layers, and 64-unit output and target layers. SARDSRN added a 144-unit feature map (SARDNET) to the SRN setup. A learning rate of 0.2 was used to train both networks, while the learning and decay rates for the S ARDNET feature map input in SARDSRN were set to 0.5 and 0.9, respectively. The neighborhood was set at 6 initially and gradually reduced to 0. These parameters were found experimentally to result in the best general performance for both parsers.

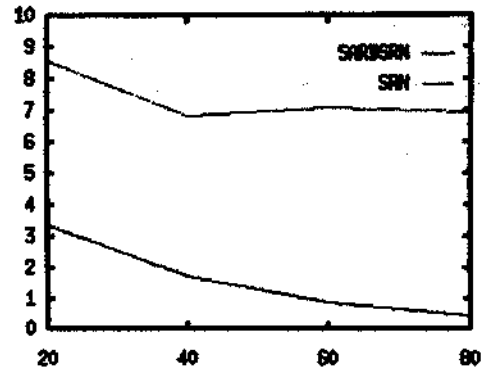

Figure 5: Results. Averages over four simulation runs using the stricter *average mismatches* per sentence measure on the test data. The SRN'S performance in all 16 runs bottomed out at a much higher error than SARDSRN, while still unable to parse all of the training sentences. SARDSRN, on the other hand, did learn to parse the training sentences, and showed very good generalization to the test sentences. These differences are statistically significant with $p < 0.0005$.

4.2 Results

The *average mismatches* performance measure reports the average number of leaf representations per sentence that are not correctly identified from the lexicon by nearest match in Euclidean distance. As an example, if the target is **[who,[liked,[the,girl]]]],** (step 11 of figure 1), but the output is **[who,[saw,[the,girl]]]],** then a mismatch would occur at the leaf labelled saw once the RAAM representation was decoded. Average mismatches provide a measure of the correctness of the information in the RAAM representation. It is a much stricter measure of the utility of the network than the standard mean squared error and was, therefore, used in our experiments.

Training took about four days on a 200 MHz Pentium Pro workstation, with SARDSRN taking about 1.5 times as long per epoch as the SRN alone. The validation error in the SRN runs quickly leveled off, and continued training did nothing to improve it On the other hand, the SARDSRN simulation runs ware still showing slight improvements when they were cut off. Figure 5 plots these performance measures averaged over the four simulation runs against the test sentences.

By all measures, SARDSRN performed significantly-even qualitatively-better than th]e standard SRN. On the training datasets, there was roughly an order of magnitude difference in both the epoch errors and the average number of mismatches per sentence between SARDSRN and SRN. These results suggest that the SRN could not even learn the training data to any useful extent, whereas SARDSRN does not appear to be nearing its limit. On the test sets, the epoch error for the SRN never fell below 0.05, and there were nearly 7 mismatches per sentence on average. Even in the most difficult case for the SARDSRN (on the 20% test dataset, in which the networks were trained on just 89 sentences, and tested on 325), these errors never reached half that level. These results show tiiat SARDSRN forms a promising starting point for parsing sentences of realistic length and complexity.

43 Example Parse

Adding SARDNET to the SRN architecture made it possible for the network to learn the parsing task. This can be shown clearly by contrasting the performances of SARDSRN and the SRN on a typical sentence, such as the one in figure 1. Neither S ARDSRN nor SRN had any trouble with the shift targets. Not surprisingly, early in training the networks would master all the shift targets in the sentence before they would get any of the reductions correct The first reduction ([the,boy] in our example) also poses no problem for either network. Nor, in general, does the second, [the,girl], because the constituent information is still fresh in memory. However, the ability of the SRN to generate the later reductions accurately degrades rapidly because the information about earlier constituents is smothered by the later steps of the parse. Interestingly, the structural information survives much longer. For example, instead of **[who,[liked,[the,girl]]]],** the SRN might produce **[who,[bit,[the,dog]]]}.** The structure of this representation is correct; what is lost are the particular instantiations of the parse tree. This is where SARDNET makes a difference. The lost constituent information remains accessible in the feature map. As a result, SARDSRN is able to capture each constituent even through the final reductions.

5 Discussion

These results demonstrate a practicable solution to the memory degradation problem of simple recurrent networks. The SRN does not have to maintain specific information about the sequence constituents, and can instead focus on what it is best at: *capturing structure.* The explicit and concise representation of the entire sequence on SARDNET also enables SARDSRN to handle long-term dependencies better than a moving-window architecture such as NARX [Mayberry and Miikkulainen, in press]. Although the sentences used in these experiments are still relatively uncomplicated, they do exhibit enough structure to suggest that much more complex sentences could be tackled with SARDSRN.

The operation of SARDSRN on the shift-reduce parsing task is a nice demonstration of holistic computation. The network is able to learn how to generate each RAAM parse representation during the course of sentence processing without ever having to decompose and recompose the constituent representations. Partial parse results can be built up incrementally into increasingly complicated structures, which suggests that training could be performed incrementally. Such a training scheme is especially attractive given that training in general is still relatively cosdy.

An extension of die SARDSRN idea, currently being investigated by our group, is an architecture where SARDNET is combined with a RAAM network. RAAM, although having many desirable properties for a purely connectionist approach to parsing, has long been a bottleneck during training. Its operation is very similar to the SRN, and it suffers from die same memory accuracy problem: with deep structures the superimposition of higher-level representations gradually obscure the traces of low-level items, and the decoding becomes inaccurate. This degradation makes it difficult to use RAAM to encode/decode parse results of realistic language. Preliminary results indicate that the explicit representation of a compressed structure formed on a SARDNET feature map, coupled with the distributed representations of the RAAM, yields an architecture able to encode richer linguistic structure.. This approach should readily lend itself to encoding the feature-value matrices used in the lexicalist, constraint-based grammar formalisms of contemporary linguistics theory, such as HPSG [Pollard and Sag, 1994], needed to handle realistic natural language.

The SARDSRN idea is not just a way to improve the performance of subsymbolic networks; it is an explicit implementation of the idea that humans can keep track of identities of elements, not just their statistical properties [Miikkulainen, 1993]. The subsymbolic networks are very good with statistical associations, but cannot distinguish between representations that have similar statistical properties. People can; whether they use a map-like representation is an open question, but we believe the SARDNET representation suggests a way to capture a lot of the resulting behavior. It is useful for building powerful subsymbolic language understanding systems, but it is also a plausible cognitive approach.

6 Conclusion

We have described an extension of the SRN called SARDSRN that combines the subsymbolic distributed properties of the SRN with the localtst properties of SARDNET. The distributed component leads to good generalization and robust cognitive properties, whereas the map provides exact representations of the sentence constituents. The results in this paper demonstrate a practicable solution to the memory degradation problem of SRNS. With SARDNET keeping track of the sequence constituents, the SRN is able to learn the structure representation necessary to perform shift-reduce parsing. This combination allows SARDSRN to learn to parse longer and more complex sentences than the SRN alone. The representative properties of SARDNET also promise to allow RAAM to encode die more complicated structures used in linguistics theory.

Acknowledgments

This research was supported in part by the Texas Higher Education Coordinating Board under grant ARP-444.
SARDSRN demo: http://www.cs.utexas.edu/users/nn/pages/-research/nlp.html.

References

[Allen, 1987] Robert B. Allen. Several studies on natural language and back-propagation. In *Proceedings of the IEEE First International Conference on Neural Networks* (San Diego, CA), volume n, pages 335-341, Piscataway, NJ, 1987. IEEE.

[Berg, 1992] George Berg. A connectionist parser with recursive sentence structure and lexical disambiguation. In *Proceedings of the Tenth National Conference on Artificial Intelligence,* pages 32-37, Cambridge, MA, 1992. MIT Press.

[Chalmers, 1990] David J. Chalmers. Syntactic transformations on distributed representations. *Connection Science,* 2:53-62,1990.

[Elman, 1990] Jeffrey L. Elman. Finding structure in time. *Cognitive Science,* 14:179-211,1990.

[Elman, 1991] Jeffrey L. Elman. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning,! A95-225,*1991.

[Hermjacob and Mooney, 1997] Ulf Hermjacob and Raymond J. Mooney. Learning parse and translation decisions from examples with rich context In *Proceedings of the 35th Annual Meeting of the ACL,* 1997.

[James and Miikkulainen, 199S] Daniel L. James and Risto Miikkulainen. SARDNET: A self-organizing feature map for sequences. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems* 7, pages 577-584, Cambridge, MA, 1995. MIT Press.

[Kohonen, 1990] Teuvo Kohonen. The self-organizing map. *Proceedings of the IEEE,* 78:1464-1480,1990.

[Kohonen, 1995] Teuvo Kohonen. *Self-Organizing Maps.* Springer, Berlin; New York, 1995.

[Mayberry and Miikkulainen, in press] Marshall R. Mayberry, m and Risto Miikkulainen. Combining maps and distributed representations for shift-reduce parsing. In Stefan Wermter and Ron Sun, editors, *Hybrid Neural Symbolic Integration.* Springer, Berlin; New York, in press.

[McClelland and Kawamoto, 1986] James L. McClelland and Alan H. Kawamoto. Mechanisms of sentence processing: Assigning roles to constituents. In James L. McClelland and David E. Rumelhart, editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 2: Psychological and Biological Models,* pages 272-325. MIT Press, Cambridge, MA, 1986.

[Miikkulainen, 1993] Risto Miikkulainen. *Subsymbolic Natural Language Processing: An Integrated Model of Scripts, Lexicon, and Memory.* MIT Press, Cambridge, MA, 1993.

[Miikkulainen, 1996] Risto Miikkulainen. Subsymbolic case-role analysis of sentences with embedded clauses. *Cognitive Science,* 20:47-73,1996.

[Miikkulainen, 1997] Risto Miikkulainen. Dyslexic and category-specific impairments in a self-organizing feature map model of the lexicon. *Brain and Language,* 59:334-366,1997.

[Munroif al, 1991] Paul Munro, Cynthia Cosic, and Mary Tabasko. A network for encoding, decoding and translating locative prepositions. *Connection Science,* 3:225-240, 1991.

[Plaut and Shallice, 1992] David C. Plaut and Tim Shallice. Perseverative and semantic influences on visual object naming errors in optic aphasia: A connectionist account. Technical Report PDP.CNS.92.1, Parallel Distributed Processing and Cognitive Neuroscience, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA, 1992.

[Plant, 1991] David C. Plaut. *Connectionist Neuropsychology: The Breakdown and Recovery of Behavior in* Lesioned *Attractor Networks.* PhD thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh, *PA,* 1991. Technical Report CMU-CS-91-185.

[Pollack, 1990] Jordan B. Pollack. Recursive distributed representations. *Artificial Intelligence,* 46:77-105,1990.

[Pollard and Sag, 1994] Carl Pollard and Ivan A. Sag. *Head-Driven Phrase Structure Grammar.* University of Chicago Press, Chicago, EL, 1994.

[Sharkey and Sharkey, 1992] Noel E. Sharkey and Amanda J. C. Sharkey. A modular design for connectionist parsing. In Marc F. J. Drossaers and Anton Nijholt, editors, *Twente Workshop on Language Technology 3: Connectionism and Natural Language Processing,* pages 87-96, Enschede, the Netherlands, 1992. Department of Computer Science, University of Twente.

[Simmons and Yu, 1991] Robert F. Simmons and Yeong-Ho Yu. The acquisition and application of context sensitive grammar for English. In *Proceedings of the 29th Annual Meeting of the ACL,* Morristown, NJ, 1991. Association for Computational Linguistics.

[Simmons and Yu, 1992] Robert F. Simmons and Yeong-Ho Yu. The acquisition and use of context dependent grammars for English. *Computational linguistics,* 18:391-418, 1992.

[St. John and McClelland, 1990] Mark F. St. John and James L. McClelland. Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence,* 46:217-258,1990.

[Stolcke, 1990] Andreas Stolcke. Learning feature-based semantics with simple recurrent networks. Technical Report TR-90-015, International Computer Science Institute, Berkeley, CA, 1990.

[Tomita, 1986] M. Tomita. *Efficient Parsing for Natural Language.* Kluwer, Dordrecht; Boston, 1986.

[Touretzky, 1991] David S. Touretzky. Connectionism and compositional semantics. In John A. Barnden and Jordan B. Pollack, editors, *High-Level Connectionist Models,* volume 1 of *Advances in Connectionist and Neural Computation Theory,* Barnden, J. A., series editor, pages 17-31. Ablex, Norwood, NJ, 1991.

[Zelle and Mooney, 1996] John M. Zelle mid Raymond J. Mooney. Comparative results on using inductive logic programming for corpus-based parser construction. In Stefan Wermter, Ellen Riloff, and Gabriela Scheler, editors, *Connectionist, Statistical, and Symbolic Approaches to Learning for Natural Language Processing,* pages 355-369. Springer, Berlin; New York, 19%.