Halmstad University Post-Print

On Learning Context-Free and Context-Sensitive Languages

Mikael Bodén and Janet Wiles

N.B.: When citing this work, cite the original article.

©2002 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Bodén M, Wiles J. On learning context-free and context-sensitive languages. New York: IEEE; IEEE Transactions on Neural Networks. 2002;13(2):491-493.

DOI: <http://dx.doi.org/10.1109/72.991436>

Copyright: IEEE

Post-Print available at: Halmstad University DiVA

<http://urn.kb.se/resolve?urn=urn:nbn:se:hh:diva-3358>

- [13] R. Palaniappan and P. Raveendran, "Single trial 40 Hz visual evoked potential detection and classification," in *Proc. IEEE 11th Workshop Statist. Signal Processing (SSP)*, Singapore, Aug. 6–8, 2001, pp. 249–252.
- [14] D. E. Rumelhart and J. L. McClelland, *Parallel Distributed Processing: Exploration in the Microstructure of Cognition*. Cambridge, MA: MIT Press, 1986, vol. 1.
- [15] J. G. Snodgrass and M. Vanderwart, "A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity," *J. Experimental Psych: Human Learning and Memory*, vol. 6, no. 2, pp. 174–215, 1980.
- [16] X. L. Zhang, H. Begleiter, B. Porjesz, W. Wang, and A. Litke, "Event related potentials during object recognition tasks," *Brain Res. Bull.*, vol. 38, no. 6, pp. 531–538, 1995.

On Learning Context-Free and Context-Sensitive Languages

Mikael Bodén and Janet Wiles

Abstract—The long short-term memory (LSTM) is not the only neural network which learns a context sensitive language. Second-order sequential cascaded networks (SCNs) are able to induce means from a finite fragment of a context-sensitive language for processing strings outside the training set. The dynamical behavior of the SCN is qualitatively distinct from that observed in LSTM networks. Differences in performance and dynamics are discussed.

Index Terms—Language, prediction, recurrent neural network (RNN).

I. INTRODUCTION

Gers and Schmidhuber [9] present a set of simulations with the so-called long short-term memory (LSTM) network on learning and generalizing to a couple of context free and a context sensitive language. The successful result is profound for at least two reasons: First, Gold [10] showed that, under certain assumptions, no super-finite languages can be learned from positive (grammatically correct) examples only. The possibilities are thus that the network (and its environment) enforces a learning bias which enables the network to capture the language, or the predictive learning task implicitly incorporates information of what is ungrammatical (cf. [14]). Second, the network establishes the necessary means for processing embedded sentences without requiring potentially infinite memory (e.g., by using stacks). Instead the network relies on the analog nature of its state space.

Contrary to what is claimed in [9], the LSTM is not the only network architecture which has been shown able to learn and generalize to a context sensitive language. Specifically, second-order sequential cascaded networks (SCNs) [12] are able to learn to process strings well outside the training set in a manner which naturally scales with the length of strings [4]. First-order simple recurrent networks (SRNs) [8] induce

TABLE I

RESULTS FOR RECURRENT NETWORKS ON THE CSL $a^n b^n c^n$, SHOWING (FROM LEFT TO RIGHT) THE NUMBER OF HIDDEN (STATE) UNITS, THE VALUES OF n USED DURING TRAINING, THE NUMBER OF SEQUENCES USED DURING TRAINING, THE NUMBER OF FOUND SOLUTIONS/TRIALS, AND THE LARGEST ACCEPTED TEST SET

Network	Hidden Units	Train. Set $[n]$	Train. Str. $[10^3]$	Sol./Tri.	Best Test $[n]$
SCN/BPTT[4]	2	1, ..., 10	max 20	23/1000	1, ..., 18
SRN/IHC[6]	3	1, ..., 12	na	na	1, ..., 12
LSTM[9]	na	1, ..., 10	avg 54	10/10	1, ..., 52

similar mechanisms to SCNs but have not been observed to generalize [6].

As reported, the LSTM network exhibit impressive generalization performance by simply adding a fixed amount to a linear counter in its state space [9]. Notably, both SRNs and SCNs process these recursive languages in a qualitatively different manner compared to LSTM. The difference in dynamics is highlighted as it provides insight into the issue of generalization but also as it may have implications for application and modeling purposes. Moreover, complementary results to Gers and Schmidhuber's [9] treatment are supplied. We focus on the SCN since it clearly demonstrates generalization beyond training data.

II. LEARNING ABILITY

Similar to [9] networks are trained using a set of strings, called S , generated from $a^n b^n c^n$ where $n \in 1, \dots, 10$. Strings from S are presented consecutively, the network is trained to predict the next letter, and n is selected randomly for each string.¹ Contrary to [9] we do not employ *start-of-string* or *end-of-string* symbols (the language is still context-sensitive). The crucial test for successfully processing a string is based on predicting the first letter of the next string.

The SCN has three input and three output units (one for each symbol; a , b and c). Two sigmoidal state units² are sufficient. Consequently, the SCN has a small and bounded state space $[0, 1]$ in contrast with the LSTM which is equipped with several specialized units of which some are unbounded and some bounded.

Backpropagation through time (BPTT) is used for training the SCN. The best SCN generalizes to all strings $n \in 1, \dots, 18$ (see Table I) and the best LSTM manages all strings $n \in 1, \dots, 52$ [9]. BPTT suffers from a "vanishing gradient" and is thus prone to miss long-term dependencies [1], [11]. This may to some extent explain the low proportion of SCNs successfully learning the language (see Table I). In a related study the low success rate and observed instability during learning are partly explained by the radical shifts in dynamics employed by the network [5]. Chalup and Blair [6] trained a three hidden unit SRN to predict $a^n b^n c^n$ using an incremental version of hill-climbing (IHC; see Table I).

III. GENERALIZATION

A. Infinite Languages

It is important to note that neither Gers and Schmidhuber [9] nor we are actually training our networks using a nonregular context sensitive language. A finite fragment of what a context sensitive grammar generates is a straightforwardly regular language—there are only ten

¹The target is only the next letter of the current string and not (as in [9]) all possible letters according to the grammar. According to [9] LSTM performs similarly in both cases.

²The logistic activation function was used.

Manuscript received September 13, 2001.

M. Bodén is with the School of Information Science, Computer and Electrical Engineering, Halmstad University, 30118 Halmstad, Sweden (e-mail: mikael.boden@ide.hh.se).

J. Wiles is with the School of Information Technology and Electrical Engineering, University of Queensland, St. Lucia 4072, Australia (e-mail: janetw@itee.uq.edu.au).

Publisher Item Identifier S 1045-9227(02)01792-7.

Fig. 1. The activation trajectory in the state space of an SCN, generated when presenting $a^8b^8c^8$, is shown as a solid line. The decision boundaries in output space (0.5) are shown as dotted lines. When the first a is presented the state quickly aligns with the fixed point of the a -system (1.00, 0.56) and starts oscillating toward it. When the first b is presented the state is attracted to the b fixed point (0.28, 0.45) but the attraction is only in one principal dimension and as oscillation continues the state is repelled from the same fixed point in the other principal dimension. The repel rate matches the a count to signal when the first c is expected. The attract rate of the b fixed point is used for determining when the final c has been presented. The c fixed point (0.36, 0.89) is repelling by oscillation. The activation crosses the decision boundaries one time step ahead of each symbol shift due to a temporal delay of the SCN.

possible strings and there is a finite automaton that will recognize each and reject all others. To test if $a^{11}b^{11}c^{11} \in S$ is to test if the network employs some other means than those expected in a finite automaton. However, networks have been observed to do so even when trained on regular languages [2]. So, to claim complete success, we need to present a proof that the network is solving all possible instances, unless we test for an infinite number of strings. Neither the LSTM nor the SCN or SRN succeed in this strict sense (admittedly, Gers and Schmidhuber test for a very large number of strings). To this end, the underlying mechanisms need to be established.

Turing machines are able to process infinitely long strings of any language since they have access to an infinite memory. Linear bounded automaton is (at most) linear in the input size and if an arbitrary long string—generated by a context sensitive grammar—is presented it can assume sufficient memory resources. Can we establish that our networks process strings in a manner which similarly scales with the size of the input?

B. Processing Mechanisms

As noted by Gers and Schmidhuber, SRNs learn simple context free languages. One of the studied languages, $a^n b^n$, was also tested on SCNs [4]. SCNs are as capable of inducing mechanisms for $a^n b^n$ as SRNs. After successful training, SCNs also exhibit similar dynamics as SRNs for processing. Interestingly—in particular in the light of the clear distinction between context free and context sensitive languages made by classical linguists—qualitatively similar dynamics is induced for $a^n b^n c^n$ [4].

For $a^n b^n$ we observe two principal types of dynamics which generalize beyond the training set [3].

- Oscillation toward an attractive fixed point while counting as and oscillation from a repelling fixed point while counting bs . The

oscillation rate is set so that the a -count matches the b -count. Only one principal dimension is required to keep track of one counter.

- Spiraling toward a fixed point while counting as and spiraling counterwise from a near fixed point. The spiraling rate is set so that the a -count and b -count match. At least two dimensions are required to implement the spiral.

A majority of the successful SCNs, in a large set of simulations, employed oscillation. The SRN only employed oscillation.

For $a^n b^n c^n$ we observe only one type of dynamics which generalizes beyond training data: oscillation toward a fixed point while counting as , dual-pronged oscillation with reference to a saddle fixed point while counting bs and oscillation from a third fixed point while counting cs (for a typical example see Fig. 1). The oscillation rates are set so that the a -count matches the b -count and the b -count matches the c -count. The second fixed point requires two principal dimensions: one in which the fixed point is repelling (matching the a -count) and one in which the fixed point is attractive (to match the c -count). By linearizing the system around the fixed points it is possible to characterize the dynamics in terms of eigenvalues and eigenvectors [13]. In [4] it was established, using linearization, that the dynamics for processing $a^n b^n$ and $a^n b^n c^n$ are qualitatively the same.

According to the diagrams presented in [9] state units in LSTM networks seem to count by increasing or decreasing activation (referred to as the “cell state” in [9]) monotonically. Basically, every a presented increases the activation of a cell, and every b decreases the activation of the same cell, using a matched rate. When the activation level is below a threshold the network will predict the symbol shift. For $a^n b^n c^n$ at least two cells are required but the same principle applies. This mechanism seems to rely on unbounded activation functions and an infinitely large state space. In SRNs and SCNs monotonic counters have also

been found—but they do not generalize [16]. However, since bounded sigmoidal activation functions are used on state units the activation quickly saturates.

IV. DISCUSSION

Monotonic counters are obviously much more stable—both in terms of learning (the LSTM network was never observed to loose track of a solution) and processing (the LSTM generalized remarkably well). However, for a monotonic counter to be generalizing well it is natural to assume that an unbounded and linear state space is required. Semifractal counters as observed in bounded state spaces of SRNs and SCNs, on the other hand, degrades quickly with lack of precision. It needs to be emphasized that so does human memory [7]. The performance profile of SRNs on recursive languages correlates well with psycholinguistic behaviors [7]. Are LSTMs—given their remarkable generalization performance—cognitively plausible?

Siegelmann has proven that a recurrent network can be used to implement a universal Turing machine [15]. The proof relies on a fractal encoding of data in state space as opposed to digits on a tape. The fractal encoding has the advantage that apart from carrying counting information it can also incorporate contents. By fixing the alphabet in advance it is possible to adjust the length of trajectories in state space into smaller or larger fractions with respect to the particular symbol and operation (push or pop). When a sequence of digits has been encoded this way—into a single value—it is possible to uniquely (and instantaneously) identify each component. The linear monotonic counter does not obviously lend itself to this added functionality as each step must be equally long. From the analysis presented in [9], a separate counter is employed by the LSTM network for each symbol and coordinated separately. The semifractal counter, on the other hand, bears many similarities to Siegelmann's proposal. It remains to be seen, however, that content-carrying oscillation can be realized and automatically induced.

The two principal approaches for scaling up with an increased input size: monotonic and fractal encoding, put different requirements on the state space. To process infinitely long strings monotonic counters require that the state space is infinitely large whereas fractal counters require a state space with infinite precision. As this brief note has only scratched the surface of the possible ways of processing recursive languages, future work should establish a precise account for the dynamics of LSTMs.

REFERENCES

- [1] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE Trans. Neural Networks*, vol. 5, pp. 157–166, Mar. 1994.
- [2] A. D. Blair and J. B. Pollack, "Analysis of dynamical recognizers," *Neural Comput.*, vol. 9, no. 5, pp. 1127–1142, 1997.
- [3] M. Bodén and A. Blair, "Learning the dynamics of embedded clauses," *Applied Intelligence: Special Issue on Natural Language Processing by Neural Networks*, 2001, to be published.
- [4] M. Bodén and J. Wiles, "Context-free and context-sensitive dynamics in recurrent neural networks," *Connection Sci.*, vol. 12, no. 3, pp. 197–210, 2000.
- [5] M. Bodén, J. Wiles, B. Tonkes, and A. Blair, "Learning to predict a context-free language: Analysis of dynamics in recurrent hidden units," in *Proc. Int. Conf. Artificial Neural Networks*, Edinburgh, U.K., 1999, pp. 359–364.
- [6] S. Chalup and A. D. Blair, "Hill climbing in recurrent neural networks for learning the $a^n b^n c^n$ language," in *Proc. 6th Int. Conf. Neural Information Processing*, Perth, 1999, pp. 508–513.
- [7] M. Christiansen and N. Chater, "Toward a connectionist model of recursion in human linguistic performance," *Cognitive Sci.*, vol. 23, pp. 157–205, 1999.
- [8] J. L. Elman, "Finding structure in time," *Cognitive Sci.*, vol. 14, pp. 179–211, 1990.
- [9] F. A. Gers and J. Schmidhuber, "LSTM recurrent networks learn simple context free and context sensitive languages," *IEEE Trans. Neural Networks*, vol. 12, pp. 1333–1340, Sept. 2001.
- [10] E. M. Gold, "Language identification in the limit," *Inform. Contr.*, vol. 16, pp. 447–474, 1967.
- [11] S. Hochreiter, "Untersuchungen zu Dynamischen Neuronalen Netzen," Diploma, Institut für Informatik, Technische Universität München, Germany, 1991.
- [12] J. B. Pollack, "The induction of dynamical recognizers," *Machine Learning*, vol. 7, p. 227, 1991.
- [13] P. Rodriguez, J. Wiles, and J. L. Elman, "A recurrent neural network that learns to count," *Connection Sci.*, vol. 11, no. 1, pp. 5–40, 1999.
- [14] D. L. T. Rohde and D. C. Plaut, "Language acquisition in the absence of explicit negative evidence: How important is starting small?," *Cognition*, vol. 72, pp. 67–109, 1999.
- [15] H. T. Siegelmann, *Neural Networks and Analog Computation: Beyond the Turing Limit*. Boston, MA: Birkhäuser, 1999.
- [16] B. Tonkes, A. Blair, and J. Wiles, "Inductive bias in context-free language learning," in *Proc. 9th Australian Conf. Neural Networks*, 1998, pp. 52–56.