A Mirror Neuron System for Syntax Acquisition

Steve Womble and Stefan Wermter

University of Sunderland, Centre of Informatics,
St. Peter's Way, Sunderland, SR6 0DD. United Kingdom
www.his.sunderland.ac.uk

Abstract. We investigate the use of a connectionist model of a mirror neuron cortical network for a context free syntax acquisition task. A finite state representation of the context free grammar is learned by an implicit knowledge system (IKS) modelled by a connectionist network. A mirror neuron system (MNS) whose evolutionary pedigree suggests adaptation for goal-directed sequential processing is used to track embedded recursions in a learned finite state model of the grammar. The mirror system modifies the output of the IKS depending on the depth of embedding. Reciprocally the IKS updates the MNS as natural 'goals' occur within a sequence during sentence production. This solves the computationally hard problem of inferring contexts from sequential input.

1 Introduction

Rizzolatti, Gallese and their colleagues have found a class of neurons in the rostral part of the ventral premotor cortex (area F5) in macaque monkeys that are active both when a monkey handles an object and when it observes an experimenter performing similar actions [6, 7, 2]. These mirror neurons are highly selective, firing only when a very specific goal-directed action is witnessed or performed. The homologue to area F5 in humans is usually taken to be Broca's area, an area traditionally associated with language processing and production [4]. Recent brain imaging experiments seem to have confirmed the existence of a similar 'human schematic matching system' [3] located in and around these areas. Consequently it has been suggested that mirror systems originally involved in goal-directed gestural actions may be involved in language processing [5].

Inspired by this work we investigate the properties of a mirror system relevant to a context free grammar acquisition task that would be difficult to learn for a conventional artificial neural network (ANN). We highlight some advantages of applying a mirror system to the acquisition task. Furthermore, we discuss the limitations of the system in a real biological environment where it would be limited by both the difficulty of learning deeper embeddings during acquisition, and by real time demands on working memory during both comprehension and production tasks.

2 The Grammar Acquisition Task

We chose for our investigation a stochastic context free noun phrase grammar. It was developed to represent the most frequent syntactic constructions in the NPL corpus, and has been investigated previously using hybrid symbolic-connectionist architectures [9].

It consists of six rule sets which themselves form two distinct sets. The noun group (NG) set consists of a noun group (NG) rule set, an adjective group (ADJG) rule set and a compound noun (CN) rule set. The noun phrase set consists of noun phrases (NP), verb phrases (VP) and prepositional phrases (PP). The construction rules for the NG set are:

NG → CN	CN → N	ADJG → ADJ
NG → DET CN	CN → N CN	ADJG → ADV ADJ
NG → ADJG CN	CN → CN ADJ CN	ADJG → ADJ ADJG
NG → DET ADJG CN		ADJG → ADJG CONJ ADJG

Grouping together these three sets yields a single self-consistent group for which a simple feed-forward ANN can accurately predict the probability distribution for the next term in a noun group by implicitly learning the conditional probability distributions and has been discussed in the context of mirror neuron systems elsewhere [10].

The rules for the noun phrase set of clauses are:

NP → NG	VP → V NP	PP → P NP
NP → NG PP	VP → V PP	PP → P VP
NP → NG VP	VP → VP CONJ VP	PP → PP CONJ PP
NP → NG CONJ NG	VP → V CONJ V NP	PP → P CONJ P NP

What makes this problem interesting is the requirement for phrase agreement either side of a conjugate in the third rule of each of the VP and PP rule sets. If both verb and prepositional phrases are active a construction which locally looks like VP CONJ PP or PP CONJ VP is possible. However if only one of the phrases is embedded (possibly multiply embedded) in a complex phrase either being produced or parsed then the grammar requires strict agreement between the conjoined phrases. The recursive nature of these noun phrases means that such agreement must be achievable over an arbitrary number of intervening elementary symbols. This makes learning with an ANN difficult due to the combinatorial complexity of the search space.

While this task may be tackled using a symbolic system, explicit statistical analysis or an ANN supported by such mechanisms [9] here we investigate how a connectionist system using only local learning rules might solve the problem in an 'online' manner given only information that is cognitively plausible.

3 The Mirror System

For the purposes of this paper our mirror system consists of two interacting parts. These are an implicit knowledge system (IKS) and a mirror neuron system

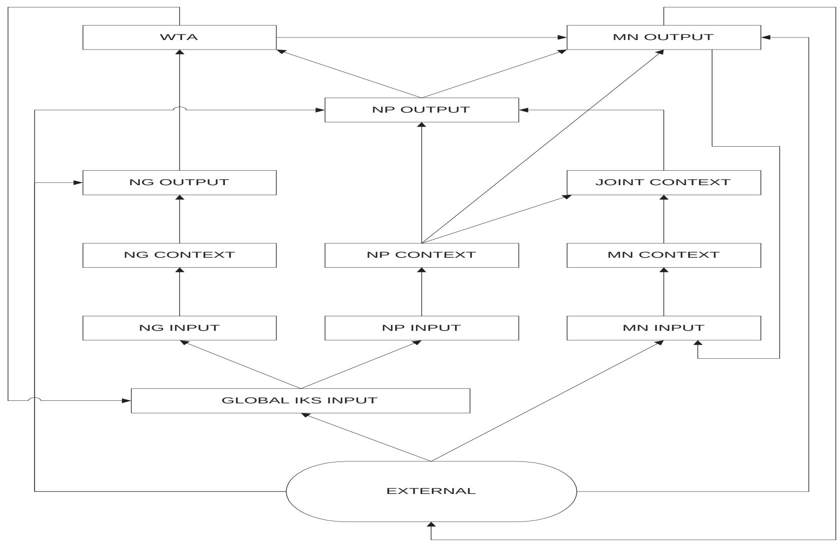

Fig. 1. High level representation of model mirror system.

(MNS). A schematic diagram of the entire system is given in Fig. 1. The IKS system is depicted on the left hand side of the diagram, and consists of the global IKS input, and the layers above it. The mirror system is depicted on the right hand side of the diagram, and consists of the MN input, and the layers above it. Connections within the system are in all cases all-to-all except for the recurrent connections from the WTA to the global IKS input and connections from the MN output to the MN input. These are 1-1 projections and are only active during production. The input layers consist of 'moving windows', containing a sparse binary encoded representation of the current symbol and a number of preceding symbols depending on the actual window width. All context layers contain nodes consisting of normalised exponentials. These learn explicit contexts found in an unsupervised manner using a model of long term heterosynaptic depression [8]. The MN context layer is used to 'clean up' the MN input representation. For the joint context layer the input is the simultaneously occurring contexts produced on each of the IKS and MNS context layers. The outputs of the IKS consist of linear nodes performing standard δ-rule error correction [1]. These learn probability distributions for the next symbol based on the context layer activity. The MN output in our model contains the neurons in our model that possess the characteristic mirror neuron properties identified by Gallese and Rizzolatti. The WTA layer performs a stochastic winner takes all and is only active when the system is used to produce sequences.

If any sequence either being produced or presented is viewed as a set of goals then the mirror neuron paradigm can easily be adapted for use. Success occurs when the resulting sequence is grammatically valid. The IKS will only allow for certain next terms to be selected or accepted as valid. However the system is now in a particular state defined by the MN input pattern, and the IKS probability distributions can then be learned to be modified to rule out incorrect sequences.

The MN input patterns can be fairly arbitrary, requiring only that they form a separable set which is isomorphic to the set of 'goals' i.e. to the relevant higher order symbols ('VP' and 'PP' are sufficient for the NPL grammar)

The mirror output layer is mapped to from the IKS output and IKS context. This layer consists of the actual mirror neurons in our model.

This characteristic output representation which after training will be active either during presentation or production contains the information that a goal, such as completing a particular phrase, has been completed. The MN output patterns are mapped onto the input patterns again using basic δ-rule error correction learning.

4 Results

As a precursor to the acquisition of the full context free grammar the system was trained using just the IKS. This allowed the system to learn the local statistics of the grammar. Once the IKS had asymptoted during training the mirror system was then activated and training continued until the system had again asymptoted. The training set consisted of 1000 strings generated stochastically from the rule sets given above. A further 10000 strings were also stochastically generated and were used as a test set for the full system.

For the mirror system there were four different input patterns; one for 'no goal', VP, PP and VP&PP combined. We used a four bit sparse binary encoding for each of these cases. For this initial investigation we dealt with deeper embeddings using a simple stack to keep track of the depth of multiple embeddings.

This is computationally convenient for the model but it is not trivial to interpret such a computational shortcut as a direct model of a biologically plausible mechanism. We discuss this issue later.

The quantitative results are given in the following table:

	IKS Only	Training Set	Test Set
Negative Log Likelihood Error	0.0644	0.0573	0.0589
One-Norm Error	0.0757	0.0690	0.0698

All figures quoted are for the average (for the given measure) over all strings tested, and are per prediction per output node. Thus the 1-norm measure gives the exact mean distance that each output node is from the correct prediction for each possible next symbol. The residual 'errors' thus include the distance from the probability distributions intrinsic to the stochastic grammar used, which for the negative log likelihood case is proportional to the entropy of the stochastic grammar.

5 Discussion and Conclusion

The reduction in error from the results for the basic IKS to the full mirror system is accounted for by the reduction in the spread of predictions for the next symbol at each point in a given sequence for which the finite state representation was unable to capture the full properties of the context free grammar. In particular with just the IKS running the system learned a mixture of probabilities for the next term in the sequence following a CONJ symbol in the constructions of conjoined VPs and PPs. The mirror system successfully disambiguated the cases for which the CONJ should be followed by a VP, a PP or either. This resulted in a sharpening of the predictions for these occurrences.

The small increase in the error between the training set and test sets highlights the advantage of our particular moving window feed-forward approach over other connectionist networks using more powerful global optimisation techniques such as stochastic gradient descent back propagation with simple recurrent networks, for which the generalisation properties of the system are dependent on the particular local optimum found during the search procedure.

Our model was successful in solving the problem of learning a complex noun phrase context free grammar from example sequences. The system can generalise probability distributions for sequence prediction over multiple complex recursions even if they exist with very low frequency in the set of training examples or are completely novel. Thus it captures the primary strength of rule based computational linguistic approaches to language acquisition using a neural system. However it also captures the cognitively not valid power of such approaches. That is, it can actually deal with arbitrarily deep layers of recursion. However the reason for this is due to our use of a stack to keep count of the depth of VP and PP recursions. This is the exact part of this system that is not yet directly related to either neuroscience evidence concerning language processing in the brain, or to psychological or cognitive models of language cognition and is the focus of our current research.

In this work we have interpreted the generation of phrase constructions as goals. The mirror system processes this goal state information and integrates it with the IKS learned local statistics of the grammar. Further the mirror system learns to identify corresponding goal states in the sequences produced by the system. At which point the activity patterns in our model mirror neurons are the same for goal states recognised either from produced or presented strings. Hence they possess the characteristic responses seen in mirror neurons in the brain. Consequently we conclude that our model mirror system provides computational evidence to support the conjecture that a similar mechanism may be involved in grammar acquisition and processing in the brain.

References

1. Christopher M. Bishop. *Neural Networks for Pattern Recognition.* Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP, United Kigdom, 1995.

2. Vittorio Gallese, Luciano Fadiga, Leonardo Fogassi, and Giacomo Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119:593–609, 1996.
3. Matthew A. Howard, F. McGlone, K. D. Singh, and N. Roberts. Brain activations to passive observation of relevant actions - an fMRI study. In Maxim Stamenov and Vittorio Gallese, editors, *Mirror Neurons and the Evolution of Brain and Language*, Advances In Consciousness Research. John Benjamins, John Benjamins, Netherlands, 2001 (In Press).
4. Ralph-Axel Muller. Innateness, autonomy, universality? neurobiological approaches to language. *Behavioural and brain Sciences*, 19:611–675, 1996.
5. Giacomo Rizzolatti and Michael A. Arbib. Language within our grasp. *Trends in the Neurosciences*, 21:188–194, 1998.
6. Giacomo Rizzolatti, R. Camarda, L. Fogassi, M. Gentilucci, G. Luppino, and M. Matelli. Functional organization of inferior area 6 in the macaque monkey: II. Area F5 and the control of distal movements. *Experimental Brain Research*, 71:491–507, 1988.
7. Giacomo Rizzolatti, Luciano Fadiga, Leonardo Fogassi, and Vittorio Gallese. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3:131–141, 1996.
8. Edmund T. Rolls and Alessandro Treves. *Neural Networks and Brain Function*, chapter 6, pages 95–135. Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP, United Kigdom, 1998.
9. Stefan Wermter. *Hybrid Connectionist Natural Language Processing*. Chapman & Hall, 2-6 Boundary Row, London., 1995.
10. Steve Womble and Stefan Wermter. Mirror neurons and feedback learning. In Maxim Stamenov and Vittorio Gallese, editors, *Mirror Neurons and the Evolution of Brain and Language*, Advances In Consciousness Research. John Benjamins, Netherlands, 2001 (In Press).