The emergence of linguistic rules in connectionist networks

Robert Kyle
r.kyle@sms.ed.ac.uk

Abstract: **In the last decade, constraint based approaches to language have proved very successful in linguistics. These approaches share a clear analog with the behaviour of neural networks, and so provide a interesting clue as to how linguistic rules might be represented in the brain. This research aims to investigate this link and examine how linguistic rules might emerge through a process of neural constraint satisfaction.**

Introduction

The question: "How are linguistic rules represented in the brain?" is a straightforward one, that has been asked many times over the last century. The fact that we still lack an adequate answer to this important question is a reflection of the shortcomings of modern Cognitive Science. Despite the overwhelming consensus that language must reside in activation of neurons in the brain, very few linguists seem willing explain how this might happen. Does the apparent difficulty of this problem tell us that we should re-evaluate many of the assumptions of modern linguistics? And do we need to reconsider our view of linguistic rules in order to better understand how they might occur as a result of neural processes?

Constraint-based approaches to language

Constraint-based approaches to language bring a fresh perspective to this important problem because they share a clear analog with patterns of activation in neural networks. This should not come as a surprise, because some of these theories (ie. Optimality Theory), take their inspiration from the constraint optimisation that occurs in networks of neurons. As with OT grammaticality judgments, neurons subjected to some input, quickly settle into the state that most satisfies the conflicting weights that constrain the network's activation. The similarity between these two processes is uncanny, because according to some, the concept of constraint satisfaction is a fundamental principle by which we can unite both high-level and neural accounts of cognition.

	Optimality Theory	Neural networks
Constraints	Universal constraint set	Weights between neurons
Constraint satisfaction	Strict domination hierarchies	Harmony maximisation
Constraint learning	The constraint demotion algorithm	Weight adjustment via Hebbian and error-driven learning
Outcome	The interaction of constraints allows for a richer description of linguistics	Neuron weights interact to determine harmonious states. The network defines a statistically rich model learned from the input

The shared properties of Optimality Theoretic and Neural Network models

Previous attempts to describe how the activity of neurons might give way to language have focused on the question of how the rules and recursion of generative linguistics might occur in the brain. Instead, Optimality theory suggests that we should focus upon how the optimisation of conflicting constraints between neurons produces the apparent rule-based behaviour we see in language.

In the eyes of many, attacking this foundational problem in Cognitive Science is an over-ambitious task for which we are not yet ready, but the new perspectives brought about by these constraint-based approaches in language suggest otherwise. The principle of constraint satisfaction provides a clear mechanism by which we can construct neural models of high-level, linguistic processes - models which are independent of the complexity of biological neurons, and rely only upon the presence of constraint-satisfaction behaviour.

A constraint-based model

Systems like Optimality Theory, which minimise multiple conflicting constraints, exhibit behaviour which can be described as Harmony maximisation. Likewise, when a network of neurons receive some input, activation rebounds through the network until a stable, self-reinforcing state is found. This settling process can also be viewed as a monotonic increase towards the most harmonious, or least conflicting state of the network. Harmony maximisation, and hence the constraint satisfaction principle, is a general property of all bi-directionally connected networks, and so high-level models are not dependent on the precise details of biological neurons.

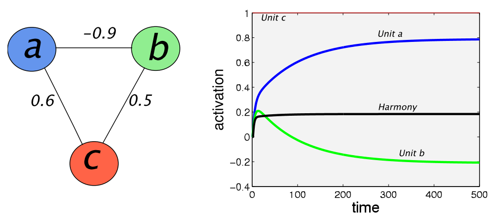

The Harmony maximisation of a simple network of neurons. Unit C is continuously active, with C=1.

In this simple example we can see that the strong inhibitory weight between unit a and b constrains the network so that both units may not be simultaneously active. We can also see that the larger weight between a and c will cause unit a to become active more quickly, and therefore win out in the inhibitory battle with b. This behaviour is captured more succinctly by the Harmony maximisation graph. The resultant, most harmonious state is a=0.79, b=-0.21 when c=1.

Now that we have seen how Harmony maximisation might provide a mechanism for constraint optimisation, how can this explain the rule-based behaviour we observe in language? The model proposes that the rules we observe in language are as the result of many thousands of interacting micro-constraints acquired as a result of linguistic experience. Harmony maximisation can be viewed as the process of disambiguating or parsing linguistic utterances, so that the final interpretation, closely obeys the statistical constraints that represent the listener's model of language. In this description, language acquisition is viewed as a process of integrative learning, over visual, auditory, social, and linguistic input - a statistical learning process which parallels the learning mechanisms that occur in networks of neurons in the brain.

Aims

The aim of the project will be to evaluate this model by analysing how well constraints learned by a neural network follow the rules described by generative linguistic theory. Ultimately the research aims to answer the question: "Do statistical approaches to linguistics offer a more rich or fundamental description of language than strictly rule-based systems?"

The method

The project will apply the model to the domain of phonology, where constraint-based approaches have had considerable success. Phonology was chosen primarily because it is one of the few linguistic domains which will not require the construction of a complex network to model learning across multiple sensory modalities. The experiment will follow that of Christiansen, Allen, and Seidenberg's (1998) model of speech segmentation. The network will be presented with a corpora of child-directed-speech, and allowed to learn the statistical relationships present in this environment.

Analysis

The performance of the model will be evaluated according to how successfully it can segment speech. The constraints that are required to capture the linguistic generalisations present in the corpus will be extracted and compared with the generalisations described by linguistic theory.

School of informatics