Learning Grammatical Structure with Echo State Networks

Matthew H. Tong, Adam D. Bickett, Eric M. Christiansen, and Gary W. Cottrell

Abstract

Echo State Networks have been shown to be effective for mathematical tasks and for modeling biological tasks such as the Mackey-Glass chaotic attractor (Jaeger, 2001). However, their performance on natural language tasks remains largely unexplored. Simple Recurrent Networks (SRN) have a long history in language modeling (Elman, 1991) and show a striking similarity in architecture to ESNs. A comparison of SRNs and ESNs on a natural language task is therefore a natural choice for experimentation. Elman applies SRNs to a standard task in statistical NLP: predicting the next word in a corpus, given the previous words. Using a simple context free grammar and an SRN with backpropagation through time (BPTT), Elman showed that the network was able to learn internal representations that were sensitive to linguistic processes that were useful for the prediction task. Here, using ESNs, we show that training of such internal representations is unnecessary to achieve levels of performance comparable to SRNs. We also compare the processing capabilities of ESNs to bigrams and trigrams. Due to some unexpected regularities of Elman's grammar, these statistical techniques are capable of maintaining longer range dependencies than might be initially expected. However, we show that the memory of ESNs in this word-prediction task extends significantly beyond that of bigrams and trigrams, enabling ESNs to make accurate predictions of verb agreement at distances over which these methods operate at chance. However, this memory is more prone to disruption by internal noise, even at small distances. The ESN performance we demonstrate might be further improved by some alterations to the standard linear fitting typically used for ESNs, as linear fits are generally not suitable to the probability-like outputs desired for a language-prediction task. Overall, our results indicate a surprising ability of ESNs to learn a grammar, suggesting that they form useful internal representations without learning them.