

Learning of Regular Languages by Recurrent Neural Networks? (Mainly Questions)

Dana Angluin

Yale University, USA

Editors: François Coste, Faissal Ouardi and Guillaume Rabusseau

Abstract

Recurrent neural network architectures were introduced over 30 years ago. From the start attention focused on their performance at learning regular languages using some variant of gradient descent. This talk reviews some of the history of that research, includes some empirical observations, and emphasizes questions to which we still seek answers.