

Exploring Small World Property in Reinforcement Learning

B. Ravindran

Indian Institute of Technology, Madras

Understanding how we are able to perform such a diverse set of complex tasks has been a central question for the Artificial Intelligence community. We hypothesize that the key to solving such tasks lies more in finding a set of sub-tasks that can easily span the set of all possible tasks, rather than finding the best sub-tasks for the goal. We model this hypothesis using the framework of reinforcement learning, and define the sub-tasks based on Kleinberg's small world model. We show some results on the small world nature of the resulting problem setting. We also demonstrate empirically that this hypothesis is valid on a general class of problems. This is joint work with Arun Chaganty and Prateek Gaur.