Evaluating the use of policy gradient optimization approach for automatic cloud resource provisioning

Włodzimierz Funika, Paweł Koperek
AGH, Faculty of Computer Science, Electronics and Telecommunication
Dept. of Computer Science, Kraków, Poland
funika@agh.edu.pl, pkoperek@gmail.com

Reinforcement learning is a very active field of research with many practical applications. Success in many cases is driven by combining it with Deep Learning. In this paper we present results of our attempt to use modern advancements in this area for automated management of resources used to host distributed software. We describe the use of three policy training algorithms from the policy gradient optimization family, to create a policy used to control the behavior of an autonomous management agent. The agent is interacting with a simulated cloud computing environment, which is processing a stream of computing jobs. We discuss and compare policy performance aspects and the feasibility to use them in real-world scenarios.

Keywords: reinforcement learning, deep neural networks, autonomous cloud management, OpenAI Gym, OpenAI Spinning Up, CloudSim Plus.