

Constraint Space Projection Reinforcement Learning



Project Manager
Dr. Riad Akrouf

Principal Investigator
Dr. Riad Akrouf

Project Term
2019 - 2020

Clusters
Lichtenberg Cluster Darmstadt

Software
PyTorch

Institute
Institute of Scientific Computing

University
Technische Universität Darmstadt

Introduction

The aim of the project is to develop a new family of optimization algorithms to handle convex constraints and evaluate them on reinforcement learning problems. In reinforcement learning, constraints are added to ensure that the learning process is safe and sound. Our optimization approach relies on finding a set of differentiable projections mapping the parameter space to a subset thereof that satisfies the constraints. The constraints and the underlying optimization routines can augment a wide range of algorithms in reinforcement learning and more broadly in machine learning.

Methods

We coded the deep learning models, the projections and the optimization routines using Pytorch that offers automatic differentiation. As we transform constrained optimization problems to unconstrained one, we are free to choose from any gradient descent algorithm. In our experiments we used the Adam optimizer that adds momentum to gradient descent, which proves critical when learning deep networks.

Results

We have shown that our optimization scheme outperforms the state-of-the-art while having linear complexity while some of the related work is quadratic in time complexity.

Discussion

We have obtained promising results towards a novel way of tackling constrained optimization in reinforcement learning. Further theoretical results in convex optimization have been recently submitted to a machine learning conference and we plan to investigate empirically the implications of the theoretical findings in a new project.

Last Update: 2023-03-16 01:29