

Project Title: Mathematical Reinforcement Learning with Emphasis on Numerical Linear Algebra

Description: This project aims to explore and advance the application of numerical linear algebra techniques within the field of reinforcement learning (RL). Reinforcement learning, a branch of machine learning, involves an agent learning to make decisions by interacting with an environment to maximize cumulative reward. The role of numerical linear algebra is critical in RL, especially in optimizing algorithms that require efficient and robust computation of solutions to linear systems, eigenvalue problems, and matrix factorizations.

Participants will focus on integrating advanced numerical linear algebra methods such as iterative solvers, Krylov subspace methods, and matrix decomposition techniques to enhance the computational efficiency of key RL algorithms like Q-learning, Policy Gradient methods, and Value Iteration. The project will involve theoretical enhancements to existing algorithms, followed by practical implementation and optimization.

A significant part of the project will be dedicated to:

1. Developing Sparse Matrix Techniques: Enhancing the scalability of RL algorithms by implementing sparse matrix operations, which are crucial for handling large state and action spaces.
2. Applying Preconditioning Methods: Investigating and developing preconditioners that can accelerate the convergence of iterative solvers used in policy evaluation and policy improvement steps.
3. Optimizing Matrix Decompositions: Utilizing techniques such as QR or Cholesky decomposition to improve the stability and efficiency of algorithms that require frequent matrix updates.

The project will also include a comparative analysis of traditional numerical methods against the newly developed techniques in various RL scenarios, assessing improvements in computational time, convergence rates, and resource utilization.

Prerequisites:

1. Strong background in numerical linear algebra and matrix computations.
2. Proficiency in programming, preferably in Python or MATLAB, with an understanding of software libraries pertinent to numerical computations (e.g., NumPy, SciPy).

3. Basic knowledge of machine learning concepts, particularly reinforcement learning, while not mandatory, would be highly beneficial.

FYP₂₀₂₄*one*

seanyshon

April 2024

1 Introduction