

Ray: A Distributed Framework for Emerging AI Applications

Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov,
Richard Liaw, Eric Liang, William Paul, Michael I. Jordan, Ion Stoica

Overview

The paper introduces a distributed framework called Ray for reinforcement learning applications.

Reinforcement learning requires many simulations which requires a large amount of compute.

Background

Not enough throughput: Map-Reduce, Apache Spark, Dryad, Dask, CIEL

Static Computation Graphs: TensorFlow, Naiad, MPI, Canary

Motivation

- High throughput, low latency.
- Dynamic graph computation.

What did they do?

Actor programming

Global Control Store

Scheduler

Fault tolerance

Ease of programming

Opinion

Limited comparison to other systems

Not exploring beyond RL