

# **Evaluation of Distributed Training with RLlib**

**Mini-Project**

**Gabriel Mahler**

# Why RLlib

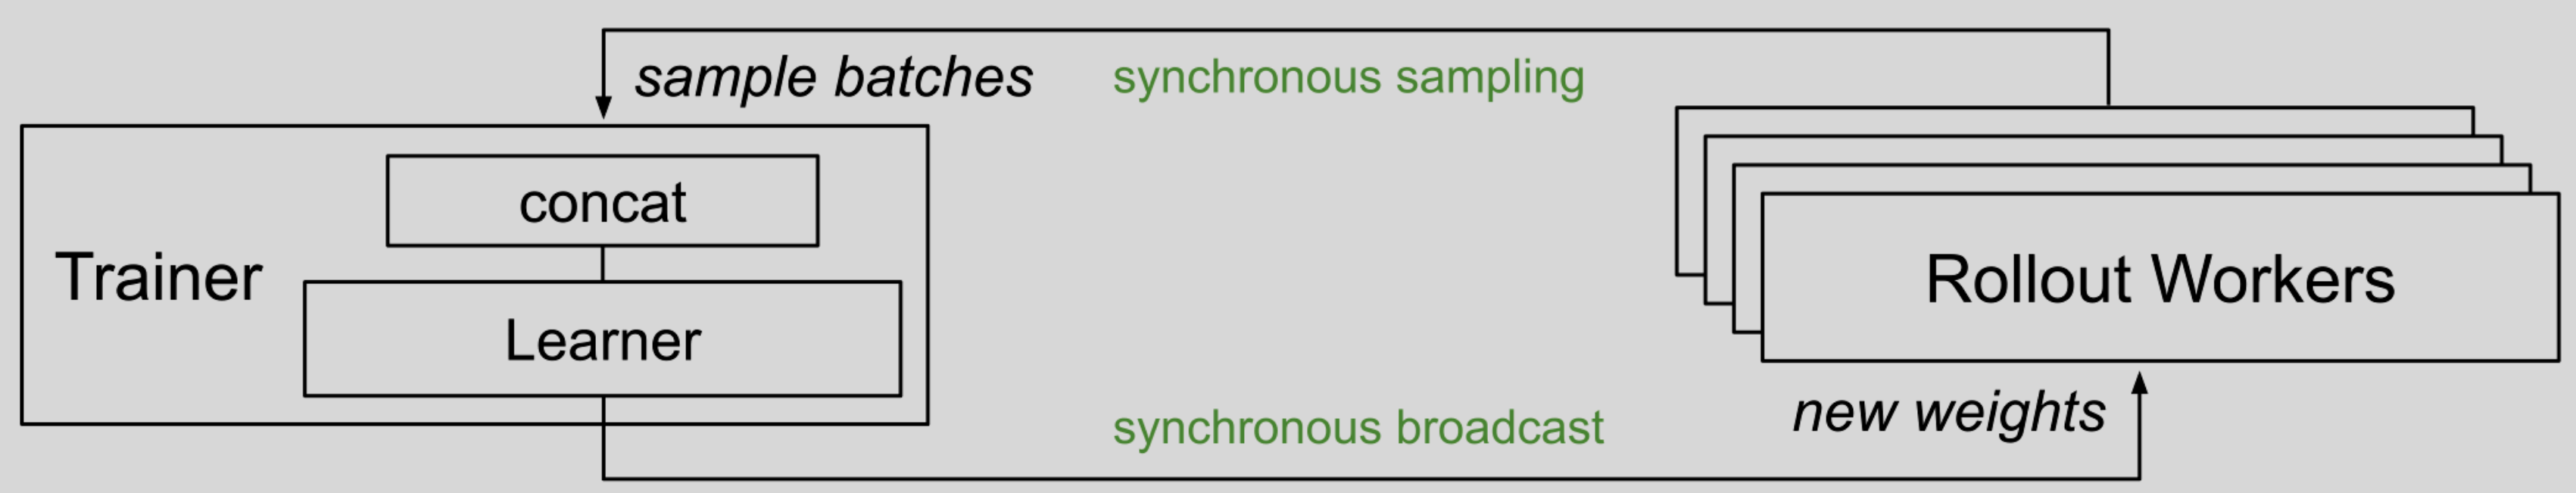
- Built on Ray
  - Ray: Framework for general-purpose distributed computing

# Why RLLib

- Built on Ray
  - Ray: Framework for general-purpose distributed computing
- Specifically focused on the training of Reinforcement Learning models

# Why RLLib

- Built on Ray
  - Ray: Framework for general-purpose distributed computing
- Specifically focused on the training of Reinforcement Learning models
- Support for distributed training



# Why RLLib

- Built on Ray
  - Ray: Framework for general-purpose distributed computing
- Specifically focused on the training of Reinforcement Learning models
- Support for distributed training
- Hierarchical and composable control

# Why RLLib

- Built on Ray
  - Ray: Framework for general-purpose distributed computing
- Specifically focused on the training of Reinforcement Learning models
- Support for distributed training
- Hierarchical and composable control
- High-level abstractions (eg. “Policy Evaluators”, “Policy Optimizers”)

# Why RLLib

- Built on Ray
  - Ray: Framework for general-purpose distributed computing
- Specifically focused on the training of Reinforcement Learning models
- Support for distributed training
- Hierarchical and composable control
- High-level abstractions (eg. “Policy Evaluators”, “Policy Optimizers”)
- Multi-agent support



# Why RLlib

- Scalable

# Why RLlib

- Scalable
- Fast

# Why RLlib

- Scalable
- Fast
- Easy to use

# Mini-Project

- Review of RLlib's distributed training

# Mini-Project

- Review of RLlib's distributed training
  - Scaling the number of devices

# Mini-Project

- Review of RLlib's distributed training
  - Scaling the number of devices
  - Parallel rollouts

# Mini-Project

- Review of RLlib's distributed training
  - Scaling the number of devices
  - Parallel rollouts
  - Model-free reinforcement learning

# Mini-Project: Why

- Haven't seen much distributed training of RL



# Mini-Project: Why

- Haven't seen much distributed training of RL
- Interesting problem

# Mini-Project: Why

- Haven't seen much distributed training of RL
- Interesting problem
- Dissertation

# Mini-Project: Evaluation

- Scaling performance (throughput, convergence)
- Experiment with several algorithms, compare performance
  - IMPALA, PPO, A3C
- Comparison to other popular frameworks (TorchRL)

# Work done so far

- Mainly research:

# Work done so far

- Mainly background research:
  - RLLib, Ray documentation

# Work done so far

- Mainly research:
  - RLLib, Ray documentation
  - “RLLib: Abstractions for Distributed Reinforcement Learning”
    - Liang et. al. (2018)

# Work done so far

- Mainly research:
  - RLLib, Ray documentation
  - “RLLib: Abstractions for Distributed Reinforcement Learning”
    - Liang et. al. (2018)
  - “RLLib Flow: Distributed Reinforcement Learning is a Dataflow Problem”
    - Liang et. al. (2021)

