RLgraph: Modular Computation Graphs for Deep Reinforcement Learning

Michael Schaarschmidt, Sven Mika, Kai Fricke, Eiko Yoneki

Reinforcement Learning (RL)

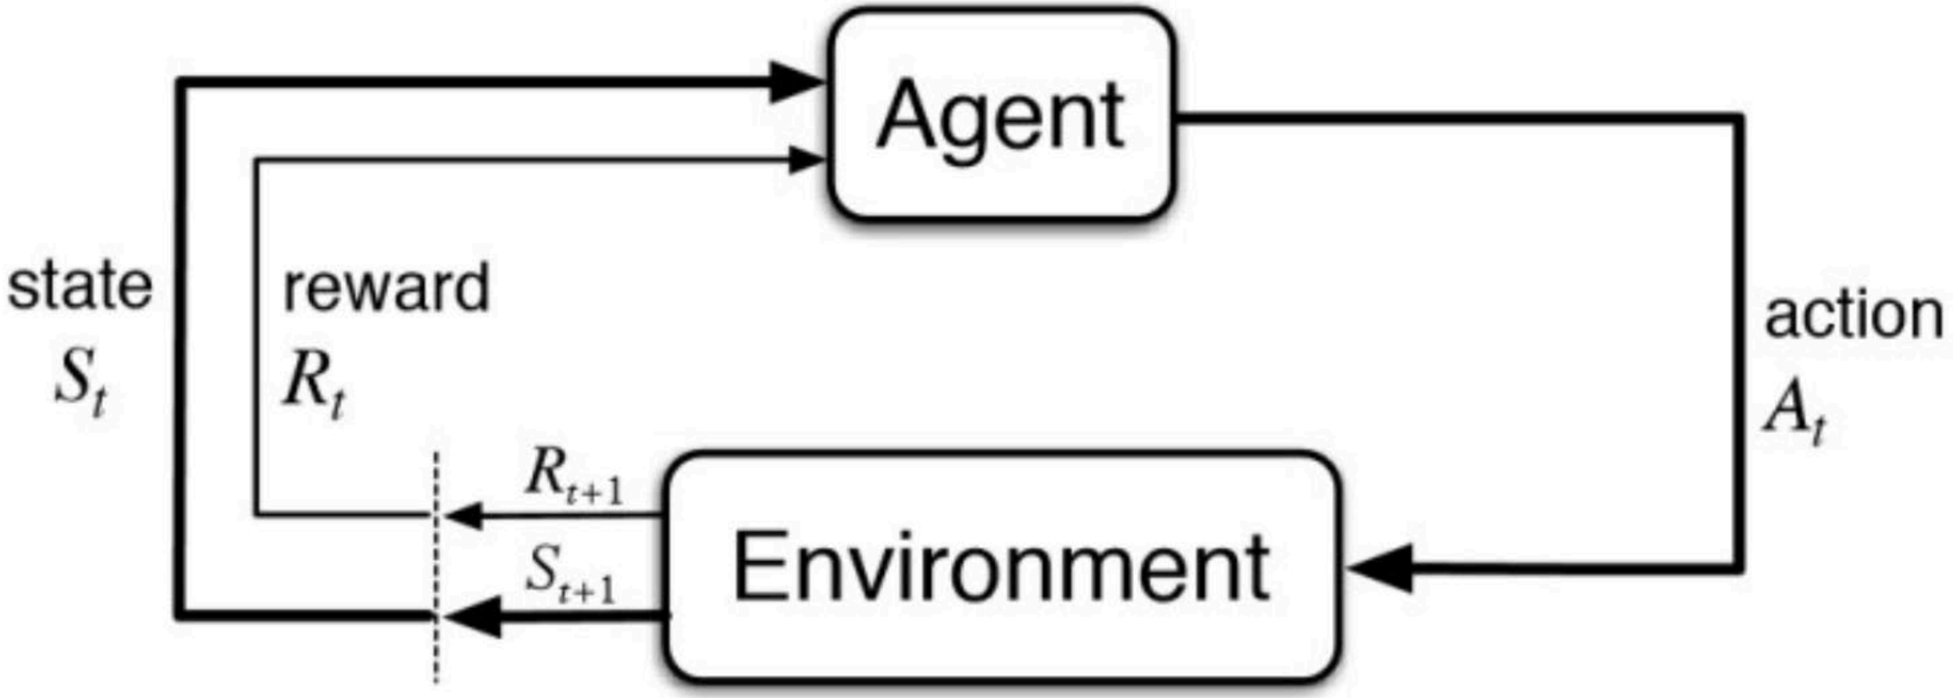

Image taken from https://www.kdnuggets.com/2018/03/5-things-reinforcement-learning.html

Supervised Learning vs RL

- Supervised Learning
  - Training data beforehand
- Reinforcement Learning
  - Learn and collect data at the same time
  - No labeled dataset
  - Sensitive to hyper parameters

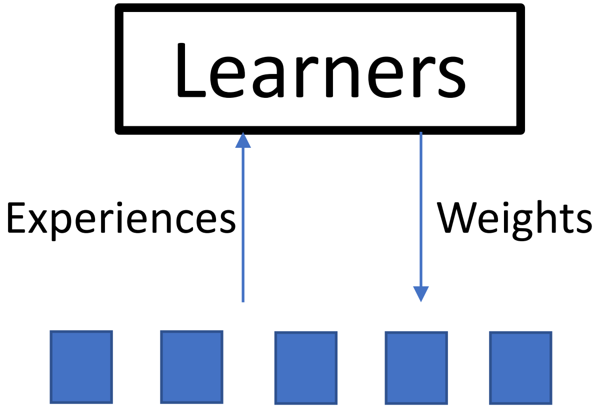

Related Works

- Existing RL Libraries
  - OpenAI baselines, TensorForce, Ray RLlib
- Pros
  - Present good results on existing environments in library
  - Code is concise
- Cons
  - Hard to adapt other environments since components tightly coupled
  - Restricted to a single backend
  - Unable to test a subcomponent individually

RLgraph

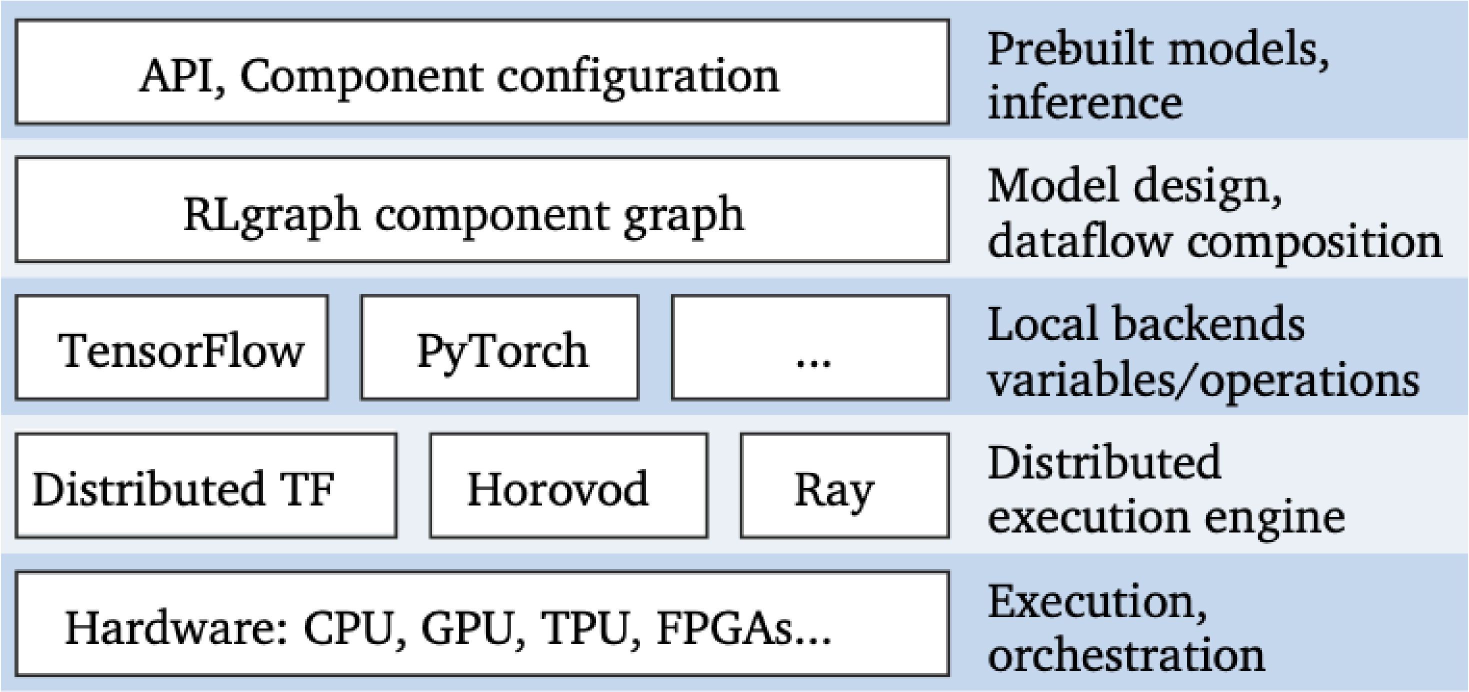

Image taken from the paper

The First Layer

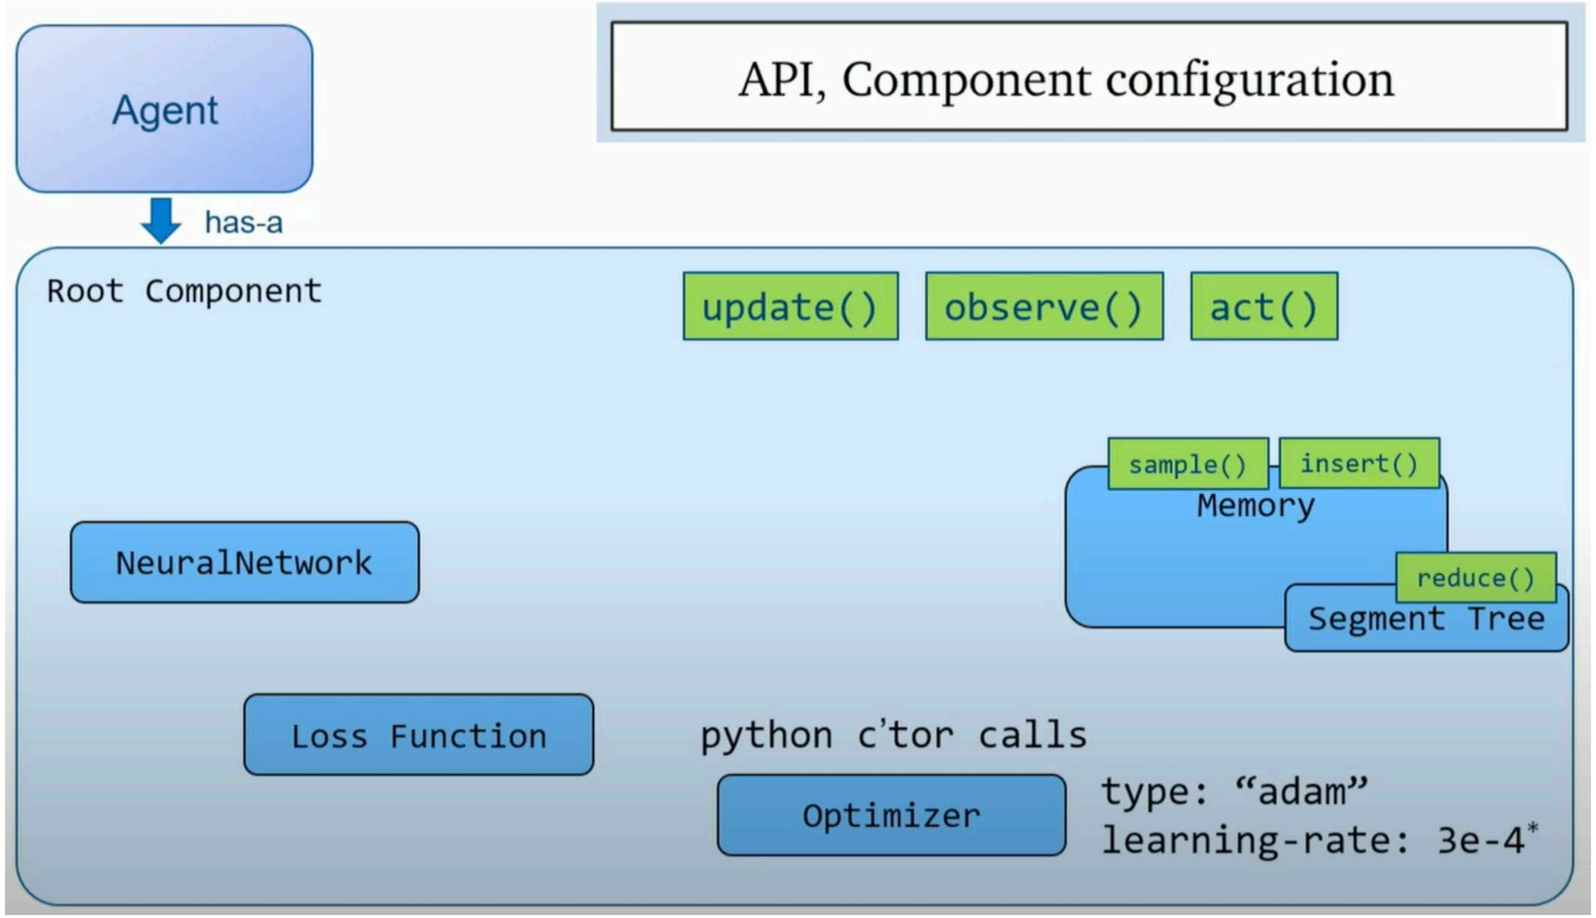

Image taken from SysML 19: https://www.youtube.com/watch?v=96cludHRSYM&t=1073s

The Second Layer

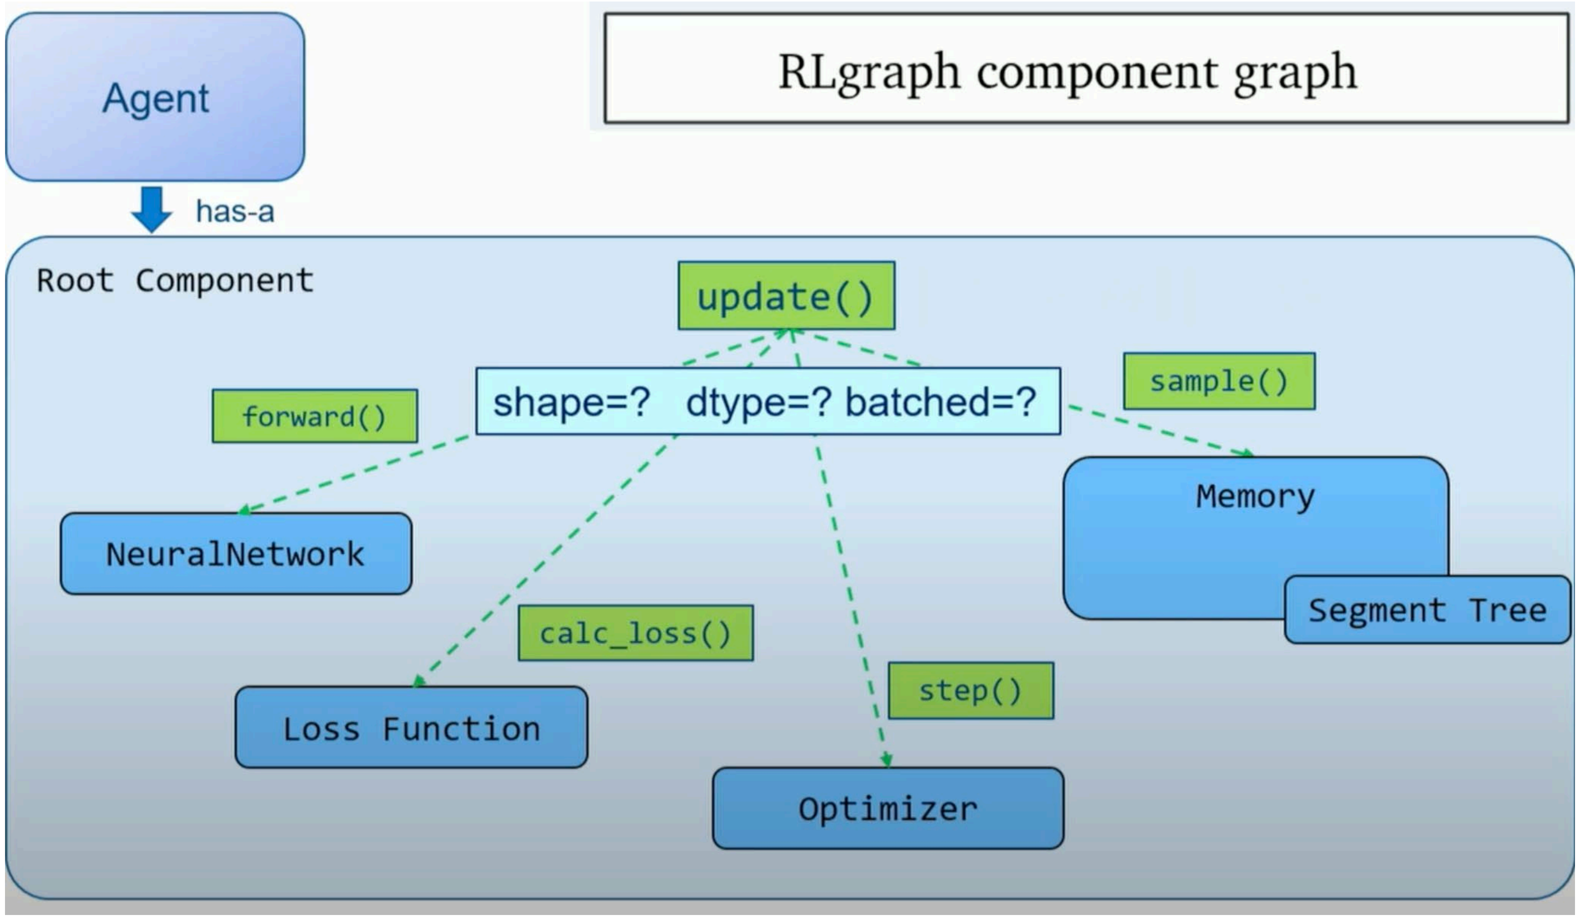

Image taken from SysML 19: https://www.youtube.com/watch?v=96cludHRSYM&t=1073s

The Third Layer

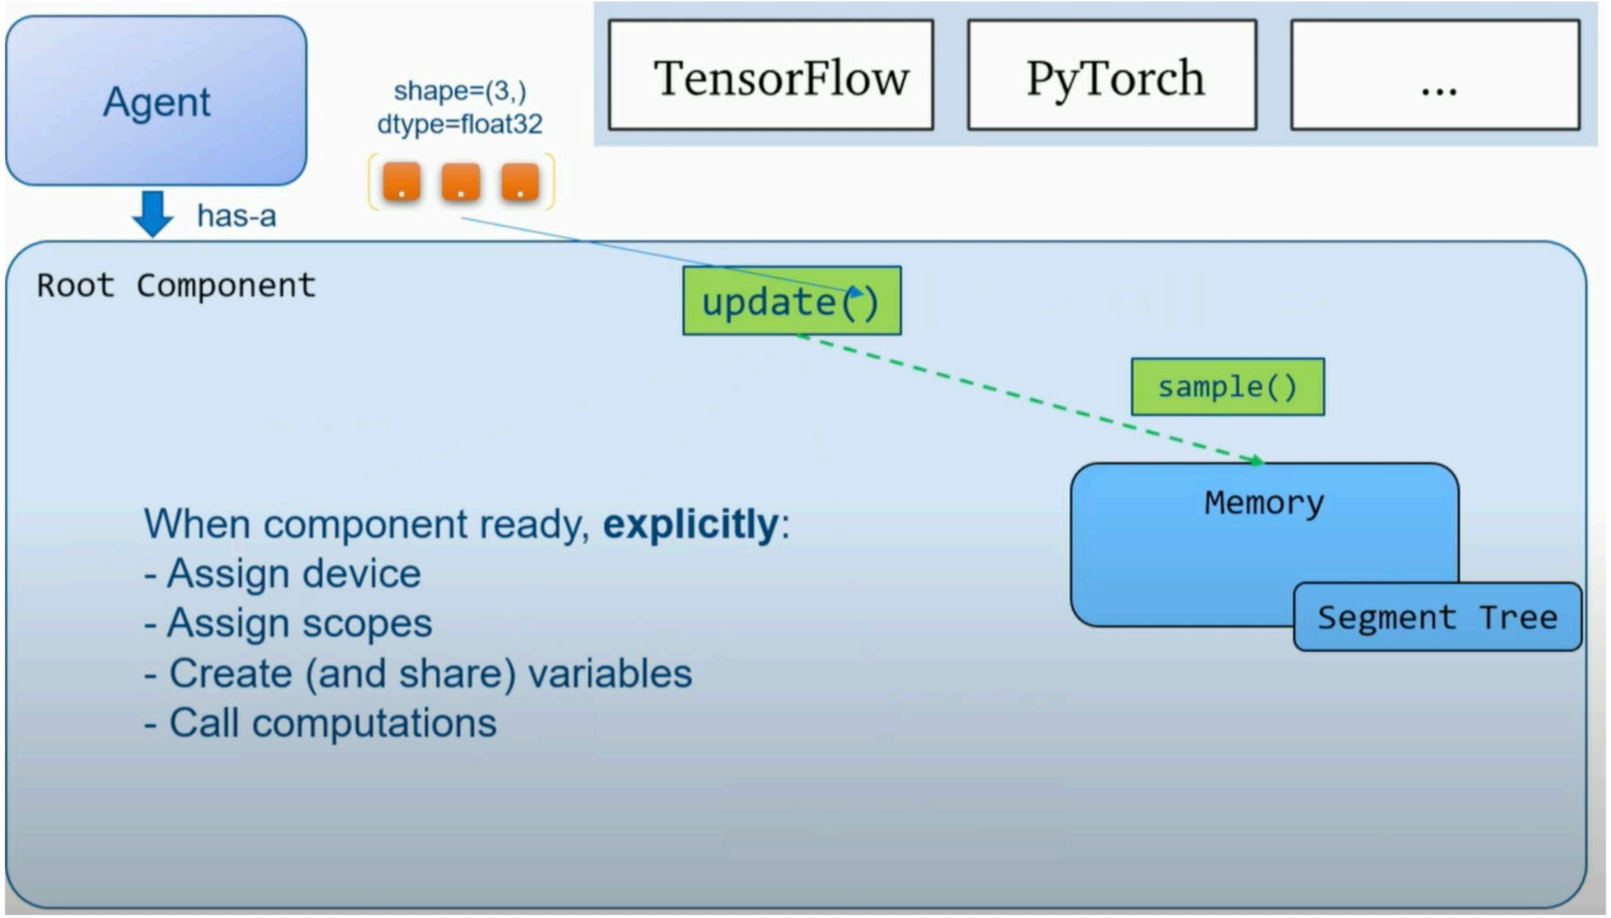

Image taken from SysML 19: https://www.youtube.com/watch?v=96cludHRSYM&t=1073s

The Fourth Layer

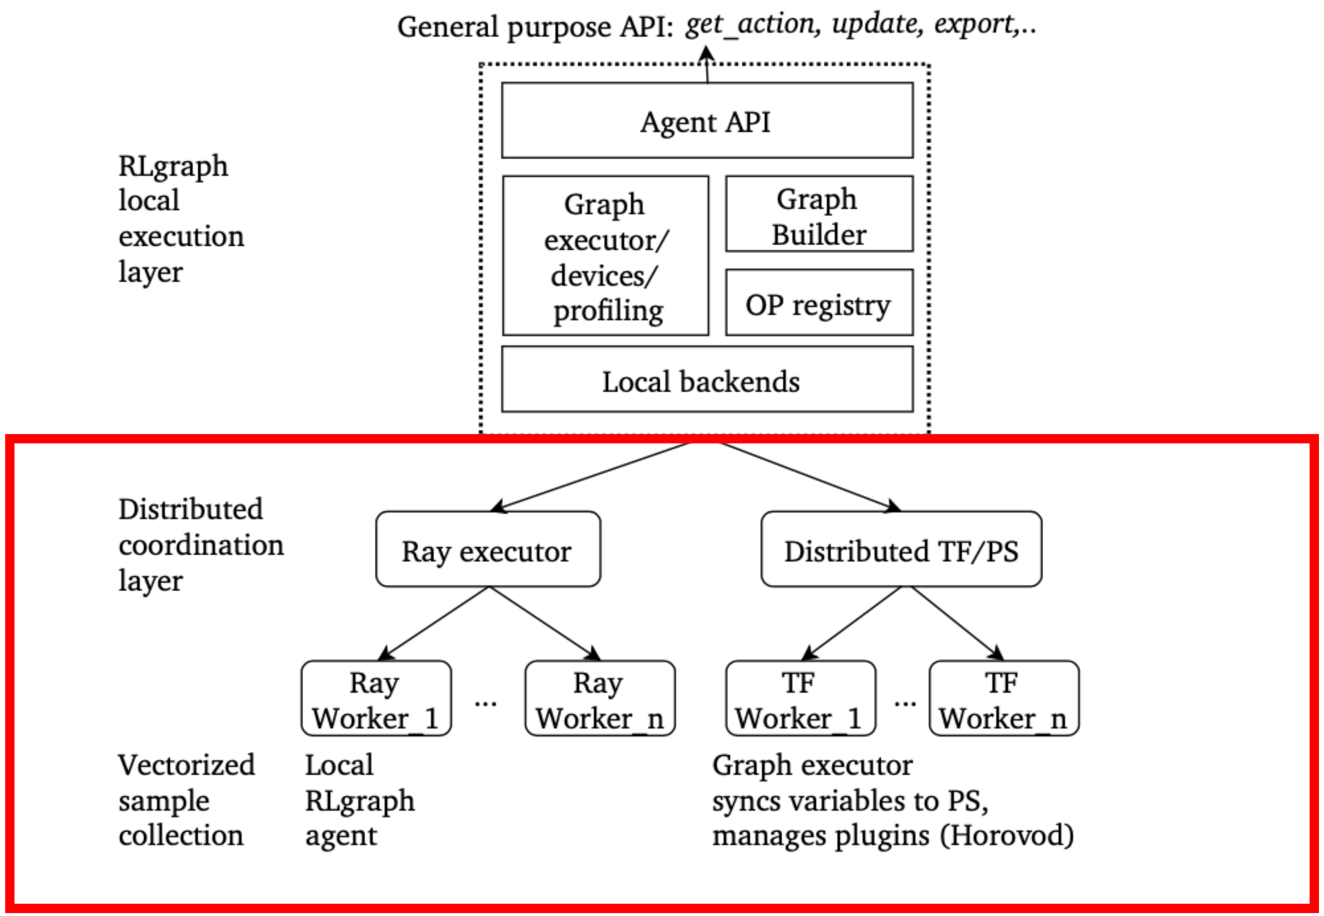

Image taken from the paper

Evaluation

Build Overhead

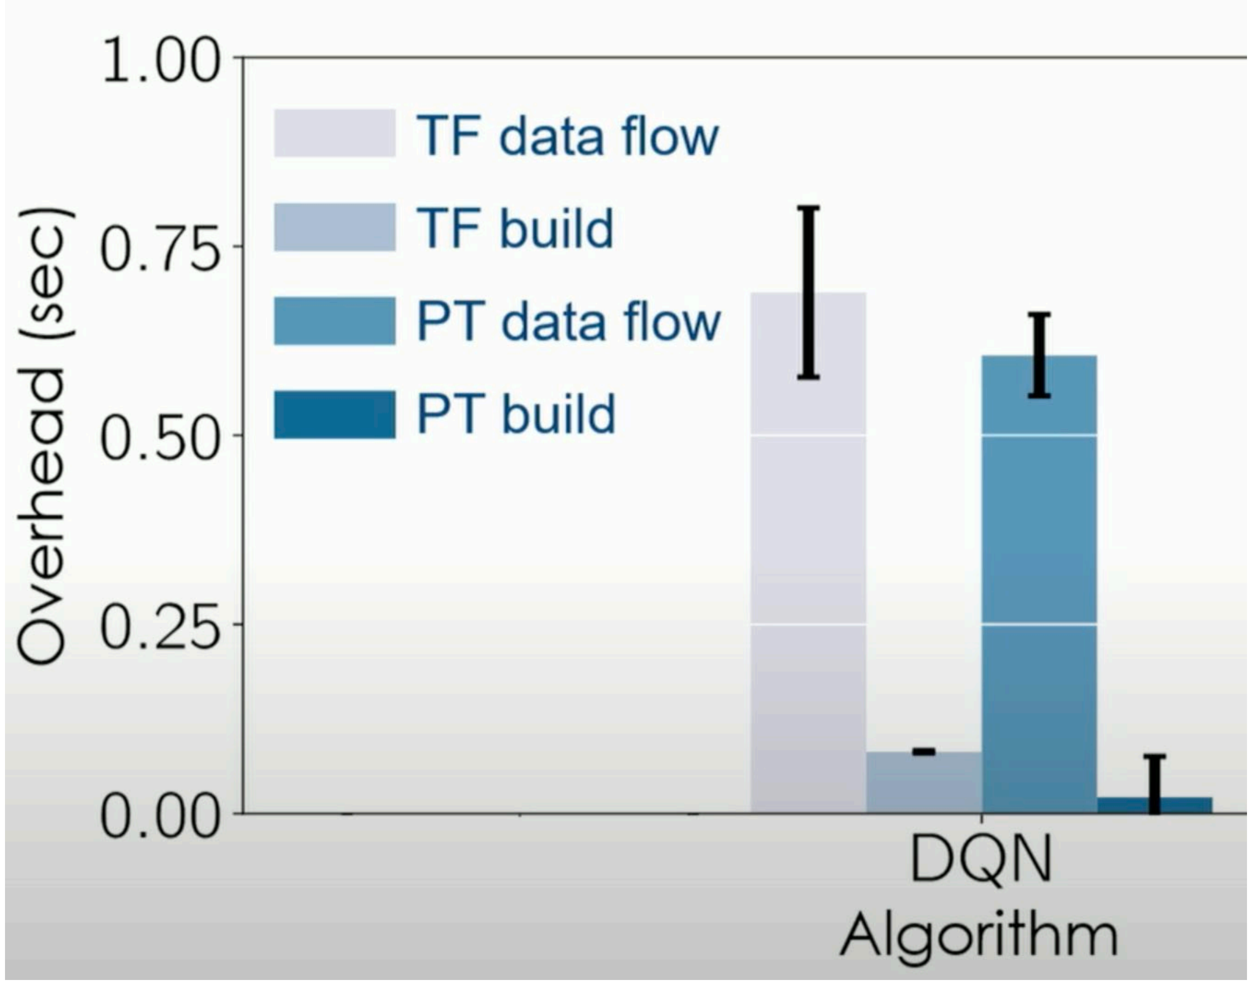

Image taken from the paper

Runtime Overhead

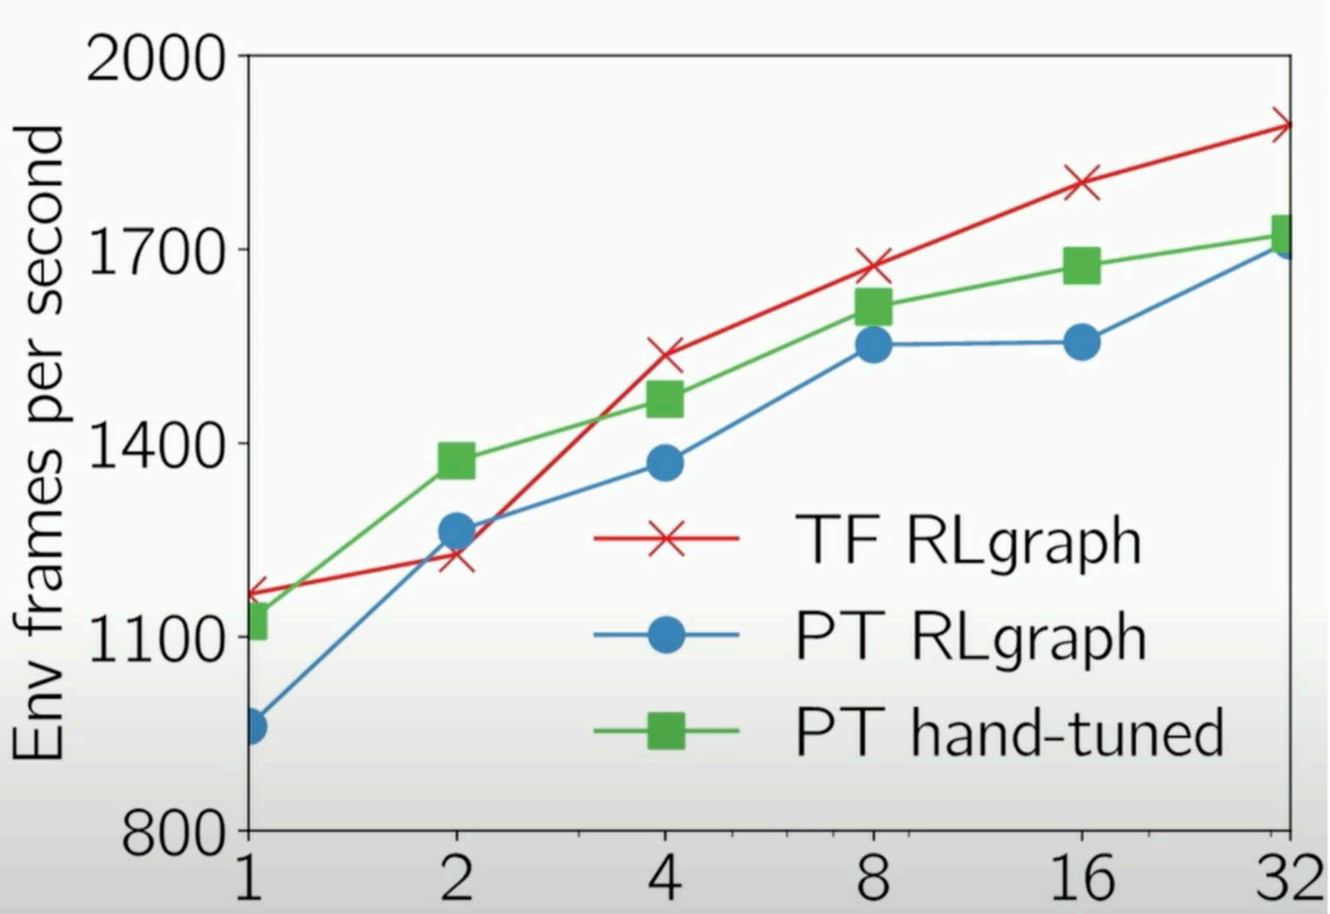

Image taken from the paper

RLgraph vs RLlib

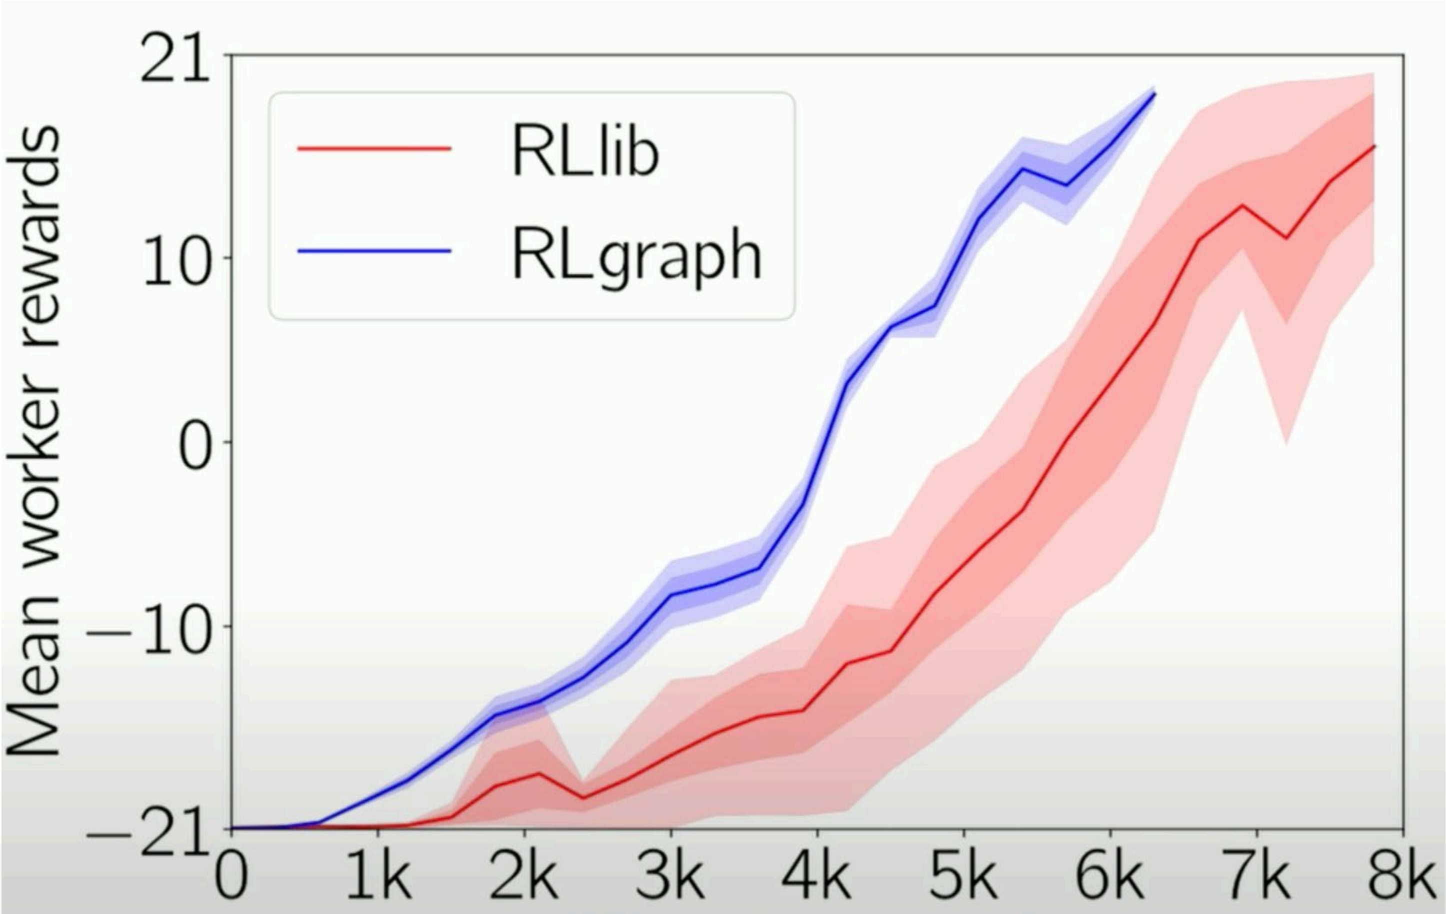

Image taken from the paper

RLgraph with multi-GPUs

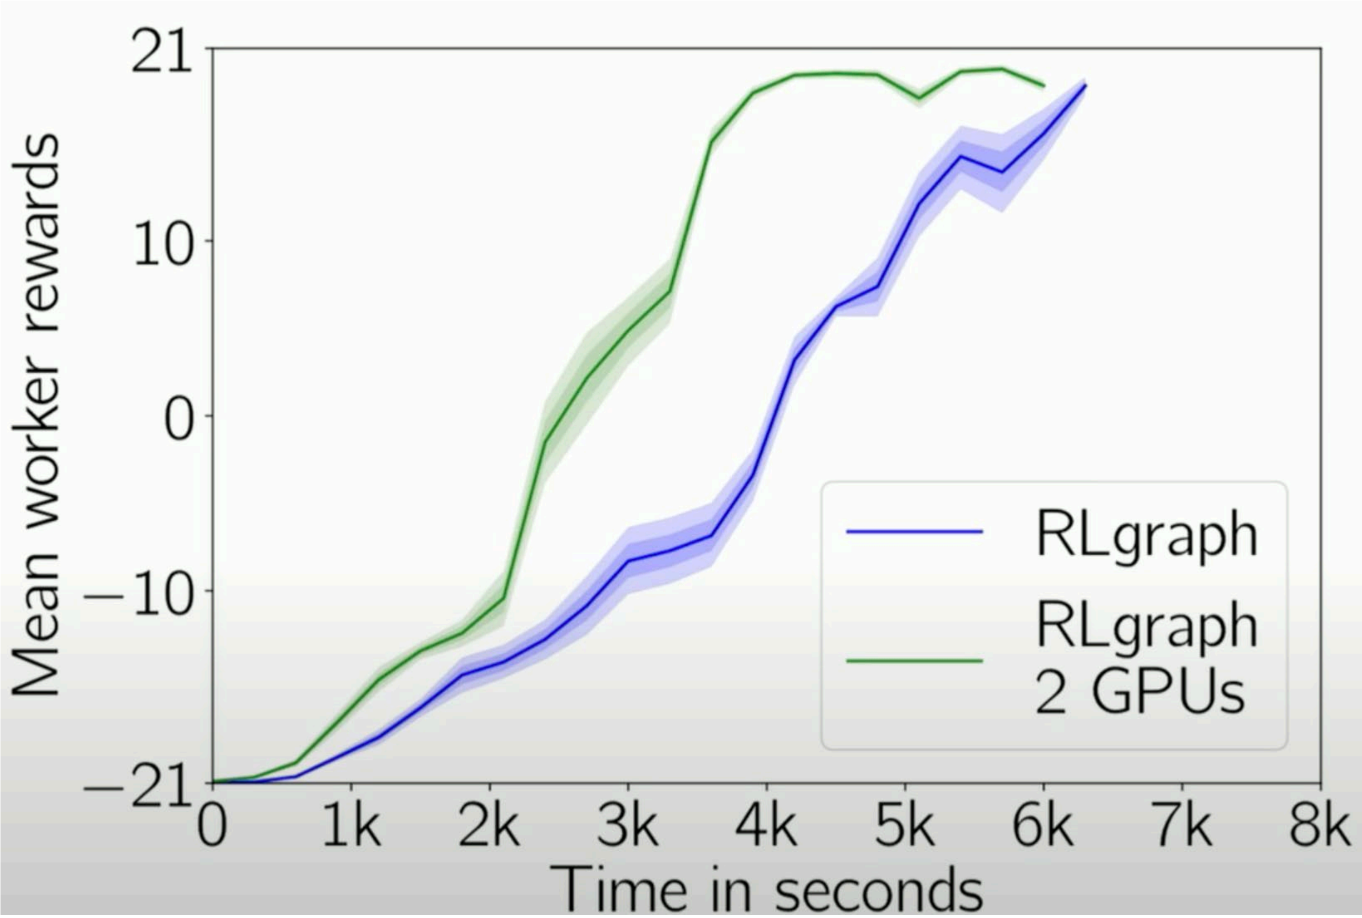

Image taken from the paper

Summary

- Introduce modularity to RL Tools
- Focus on dataflow design instead of backend tools
- Future work
  - Integrate AutoGraph / JIT Tracing into build process
- Reference
  - M. Schaarschmidt, S. Mika, K. Fricke, E. Yoneki: RLgraph: Flexible Computation Graphs for Deep Reinforcement Learning, SysML, 2019.
  - https://www.youtube.com/watch?v=96cludHRSYM&t=1073s