UNIVERSITY OF CAMBRIDGE

RLgraph: Modular Computation Graphs for Deep Reinforcement Learning

M. Schaarschmidt, S. Mika, K. Fricke, E. Yoneki at SysML, 2019

R244 Large-scale data processing and optimisation
Presentation by Martin Graf on 19/10/2022

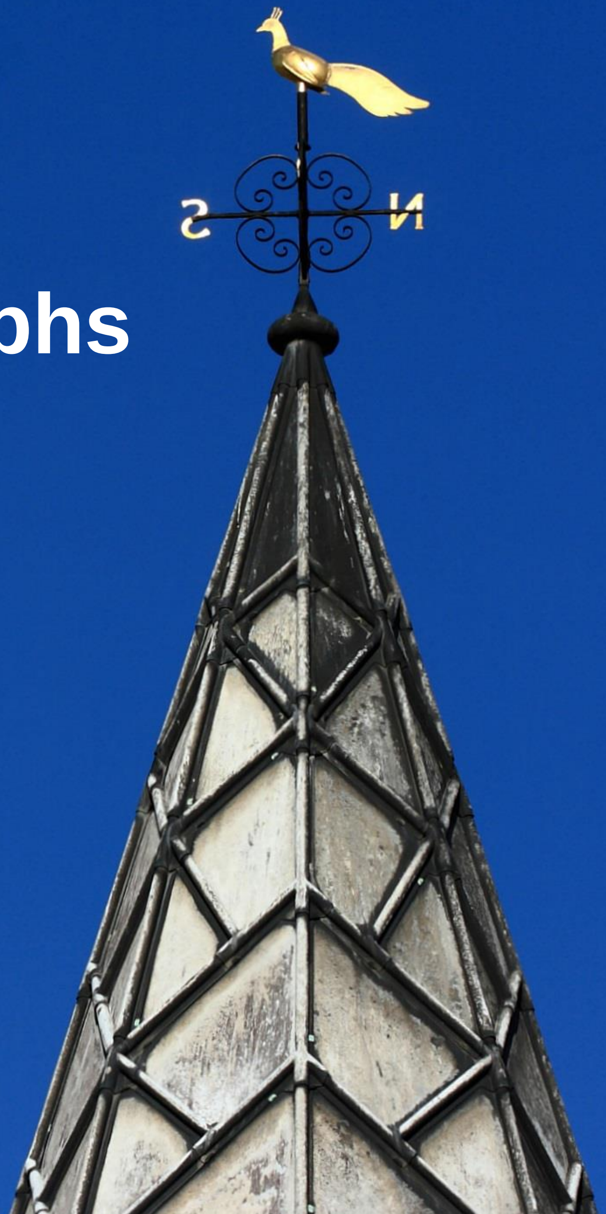

Reinforcement Learning is hard

UNIVERSITY OF CAMBRIDGE

Reinforcement Learning is hard

- Algorithmic instability

Reinforcement Learning is hard

- Algorithmic instability
- Diversity of models and optimization strategies

Reinforcement Learning is hard

- Algorithmic instability
- Diversity of models and optimization strategies
- Highly varied resource requirements

UNIVERSITY OF CAMBRIDGE

Reinforcement Learning is hard

- Algorithmic instability
- Diversity of models and optimization strategies
- Highly varied resource requirements
- Heterogeneous distributed communication patterns

Existing Reinforcement Learning Tooling

Existing Reinforcement Learning Tooling

- Reference implementations on benchmark tasks
  - OpenAI baselines (Sidor & Schulman, 2017)
  - Keras-rl (Plappert, 2016)
  - …

Existing Reinforcement Learning Tooling

- Reference implementations on benchmark tasks
  - OpenAI baselines (Sidor & Schulman, 2017)
  - Keras-rl (Plappert, 2016)
  - …
- Ray RLlib (Liang et al., 2018)

Existing Reinforcement Learning Tooling

- Reference implementations on benchmark tasks
  - OpenAI baselines (Sidor & Schulman, 2017)
  - Keras-rl (Plappert, 2016)
  - …
- Ray RLlib (Liang et al., 2018)
- TensorForce (Schaarschmidt et al., 2018)

UNIVERSITY OF CAMBRIDGE

Existing Reinforcement Learning Tooling

- Reference implementations on benchmark tasks
  - OpenAI baselines (Sidor & Schulman, 2017)
  - Keras-rl (Plappert, 2016)
  - …
- Ray RLlib (Liang et al., 2018)
- TensorForce (Schaarschmidt et al., 2018)

→ Problem: No separation of concerns

UNIVERSITY OF CAMBRIDGE

Trends in Machine Learning Tooling

UNIVERSITY OF CAMBRIDGE

Trends in Machine Learning Tooling

- Towards higher level APIs and standardization
  - Keras (Chollet et al., 2015)
  - ONNX (Facebook Inc., 2017)

UNIVERSITY OF CAMBRIDGE

Trends in Machine Learning Tooling

- Towards higher level APIs and standardization
  - Keras (Chollet et al., 2015)
  - ONNX (Facebook Inc., 2017)
- Towards better performance
  - Hardware improvements
  - Software improvements
    - Weld (Palkar et al., 2017)
    - FlexFlow (Jia et al., 2018)
    - …

UNIVERSITY OF CAMBRIDGE

Value Hypothesis

Trend: Standardization, higher level APIs

Current RL tooling: No separation of concerns

Trend: Continuous Performance Improvements

?

Value Hypothesis

Trend: Standardization, higher level APIs

Current RL tooling: No separation of concerns

Trend: Continuous Performance Improvements

RLgraph

Key Ideas

Key Ideas

- Separate execution details and user code

UNIVERSITY OF CAMBRIDGE

Key Ideas

- Separate execution details and user code
- No-code distributed computation

Key Ideas

- Separate execution details and user code
- No-code distributed computation
- Backend agnostic, high level API

UNIVERSITY OF CAMBRIDGE

Key Ideas

- Separate execution details and user code
- No-code distributed computation
- Backend agnostic, high level API
- Testable

Introducing: RLgraph

high-level **backend-agnostic** **scalable** **graph-based** **testable** library with a **component-based** **modular** build-system for designing and executing **fast**, **robust**, **incrementally testable**, and **easy to extend** or **re-use** reinforcement learning algorithms

UNIVERSITY OF CAMBRIDGE

Introducing: RLgraph

high-level backend-agnostic scalable graph-based testable library with a **component-based modular** build-system for designing and executing **fast**, **robust**, **incrementally testable**, and **easy to extend** or **re-use** reinforcement learning algorithms

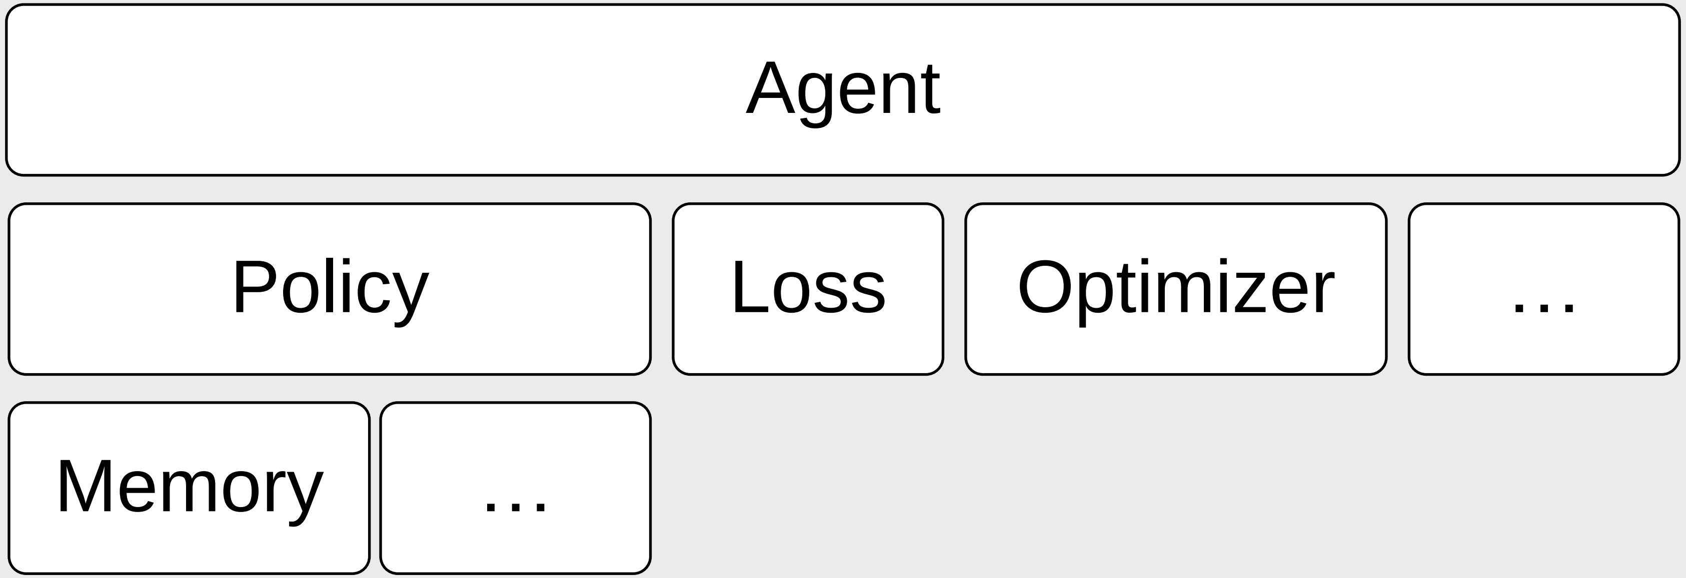

UNIVERSITY OF CAMBRIDGE

Introducing: RLgraph

high-level **backend-agnostic** **scalable** **graph-based** **testable** library with a **component-based modular** build-system for designing and executing **fast**, **robust**, **incrementally testable**, and **easy to extend** or **re-use** reinforcement learning algorithms

Multi-framework
- TensorFlow
- PyTorch

Multi-paradigm
- distributed TensorFlow (Abadi et al., 2016)
- Ray (Moritz et al., 2017)

UNIVERSITY OF CAMBRIDGE

Introducing: RLgraph

high-level backend-agnostic scalable graph-based testable library with a **component-based modular** build-system for designing and executing **fast**, **robust**, **incrementally testable**, and **easy to extend** or **re-use** reinforcement learning algorithms

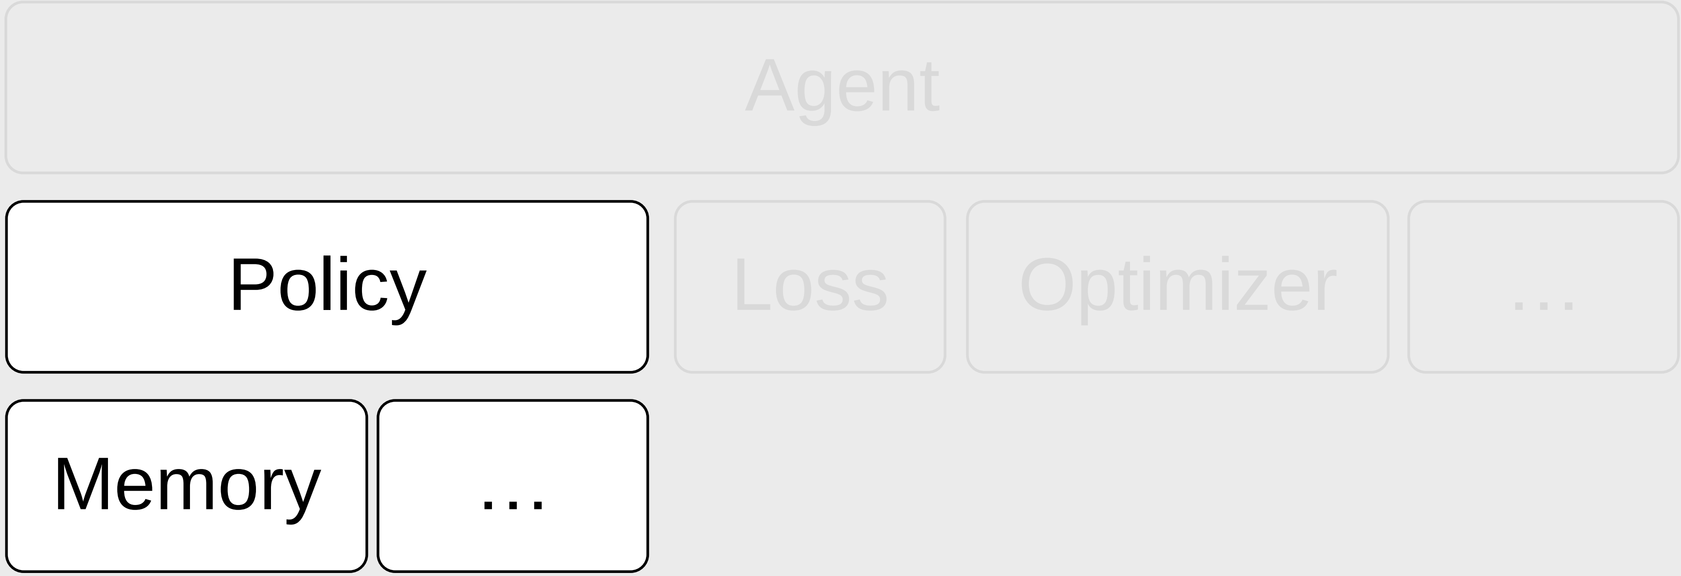

UNIVERSITY OF CAMBRIDGE

Introducing: RLgraph

high-level backend-agnostic scalable graph-based **testable** library with a **component-based modular** build-system for designing and executing fast, **robust**, **incrementally testable**, and **easy to extend** or **re-use** reinforcement learning algorithms

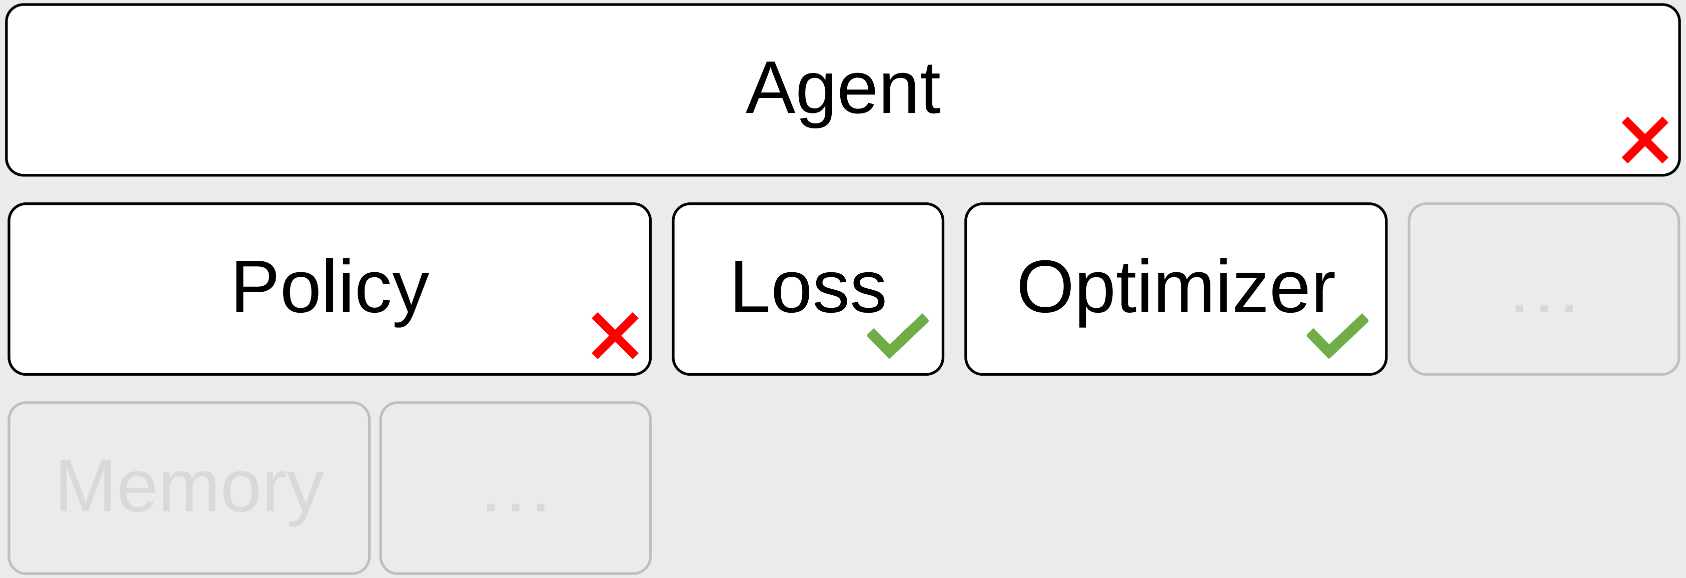

Introducing: RLgraph

high-level **backend-agnostic scalable graph-based testable** library with a **component-based modular** build-system for designing and executing **fast**, **robust**, **incrementally testable**, and **easy to extend** or **re-use** reinforcement learning algorithms

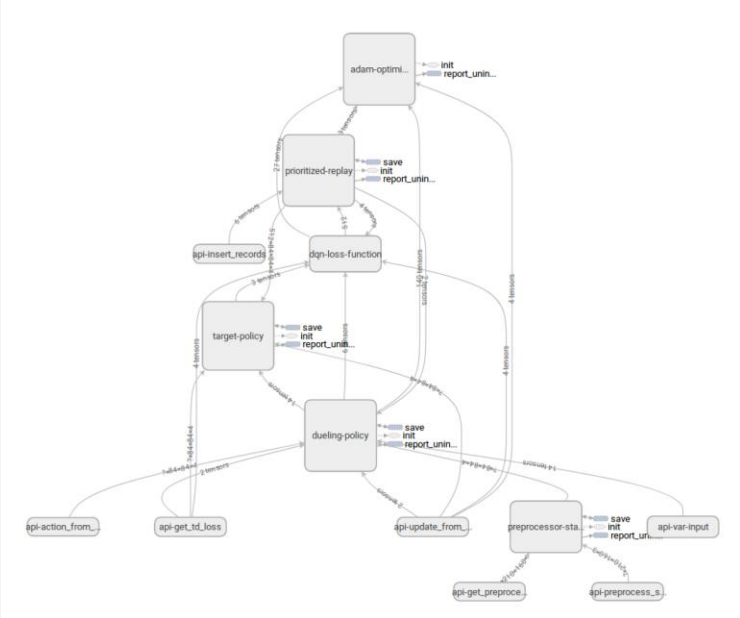

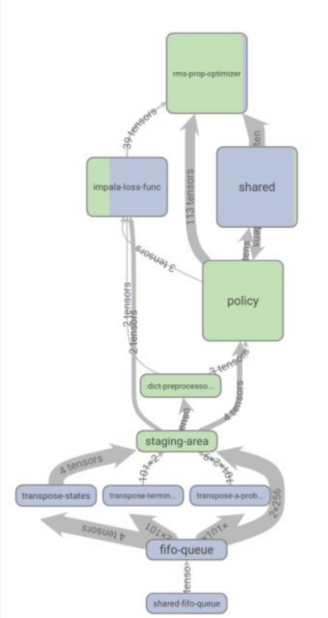

UNIVERSITY OF CAMBRIDGE

Source: M. Schaarschmidt, S. Mika, et al.: RLgraph: Modular Computation Graphs for Deep Reinforcement Learning, SysML, 2019

Introducing: RLgraph

high-level **backend-agnostic** **scalable** **graph-based** **testable** library with a **component-based** **modular** build-system for designing and executing **fast**, **robust**, **incrementally testable**, and **easy to extend** or **re-use** reinforcement learning algorithms

UNIVERSITY OF CAMBRIDGE

Performance

UNIVERSITY OF
CAMBRIDGE

Performance

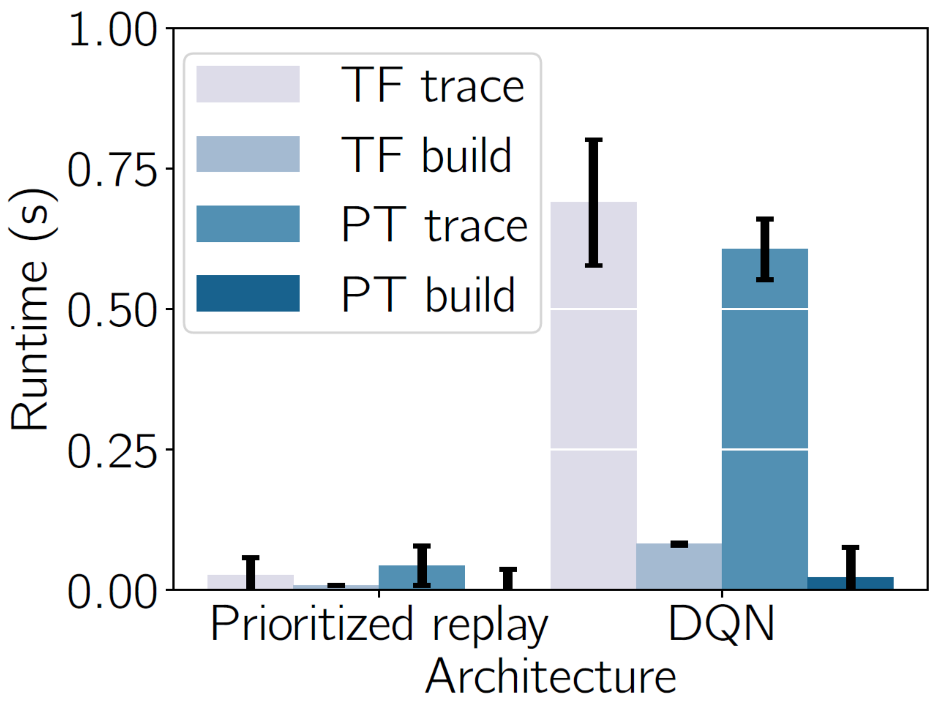

(a) Build overheads.

Source: M. Schaarschmidt, S. Mika, et al.: RLgraph: Modular Computation Graphs for Deep Reinforcement Learning, SysML, 2019

UNIVERSITY OF CAMBRIDGE

Performance

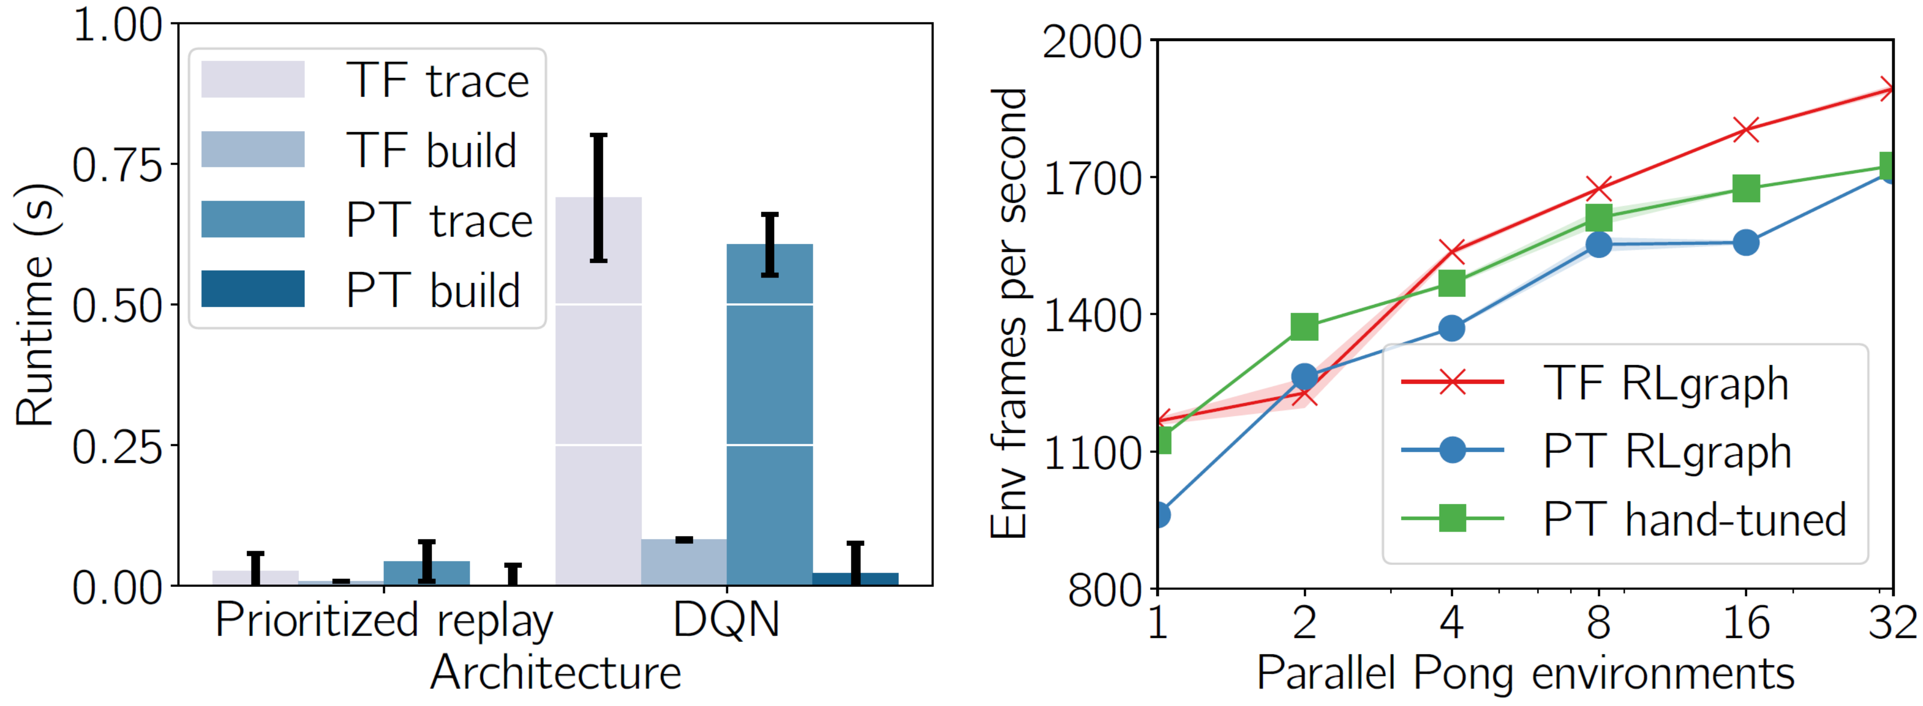

(a) Build overheads. (b) Worker act performance.

Source: M. Schaarschmidt, S. Mika, et al.: RLgraph: Modular Computation Graphs for Deep Reinforcement Learning, SysML, 2019

UNIVERSITY OF CAMBRIDGE

Performance

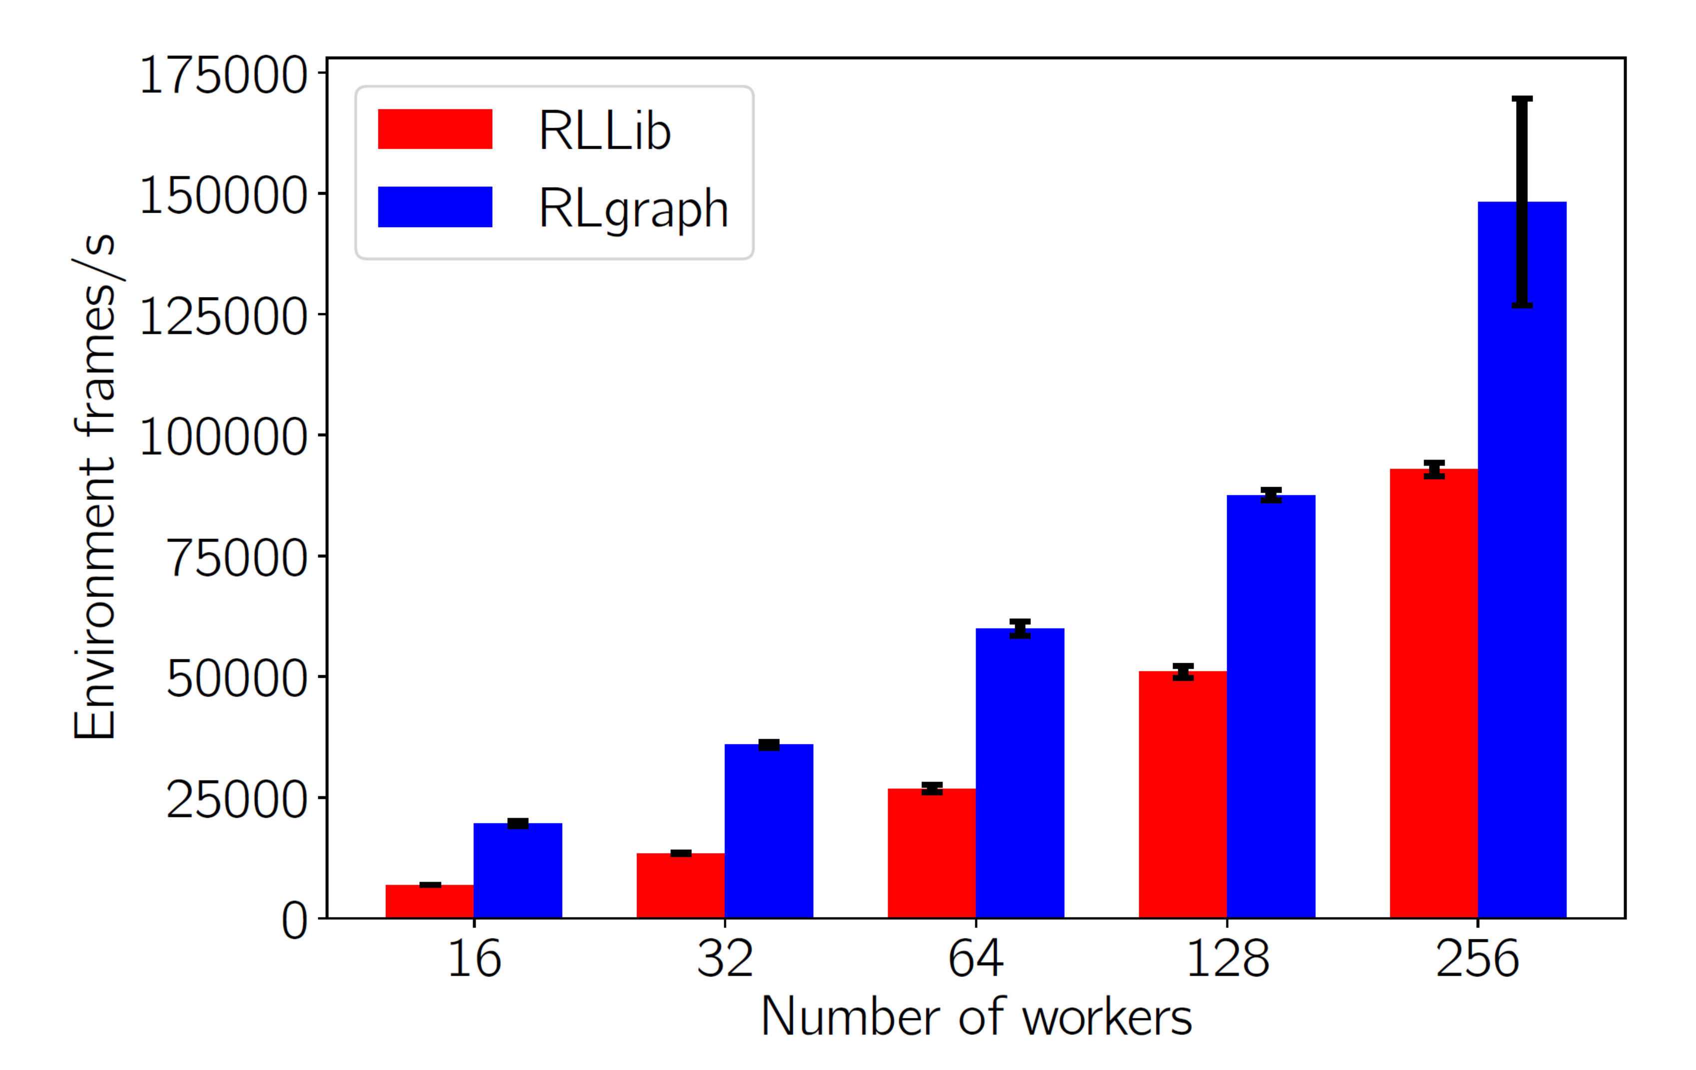

UNIVERSITY OF CAMBRIDGE

Source: *M. Schaarschmidt, S. Mika, et al.: RLgraph: Modular Computation Graphs for Deep Reinforcement Learning, SysML, 2019*

RLgraph: Modular Computation Graphs for Deep Reinforcement Learning

M Schaarschmidt, S Mika, K Fricke… - … of Machine Learning …, 2019 - proceedings.mlsys.org

Reinforcement learning (RL) tasks are challenging to implement, execute and test due to algorithmic instability, hyper-parameter sensitivity, and heterogeneous distributed communication patterns. We argue for the separation of logical component composition, backend graph definition, and distributed execution. To this end, we introduce RLgraph, a library for designing and executing reinforcement learning tasks in both static graph and define-by-run paradigms. The resulting implementations are robust, incrementally testable …

☆ Save   Cite   Cited by 13   Related articles   All 10 versions

Source: https://scholar.google.com, accessed on 18/10/2022

rlgraph / **rlgraph**   Public

RLgraph: Modular computation graphs for deep reinforcement learning

Apache-2.0 license

☆ **297** stars   **40** forks

Source: https://github.com/rlgraph/rlgraph, accessed on 18/10/2022

UNIVERSITY OF CAMBRIDGE

RLgraph: Modular Computation Graphs for Deep Reinforcement Learning

M Schaarschmidt, S Mika, K Fricke… - … of Machine Learning …, 2019 - proceedings.mlsys.org

Reinforcement learning (RL) tasks are challenging to implement, execute and test due to algorithmic instability, hyper-parameter sensitivity, and heterogeneous distributed communication patterns. We argue for the separation of logical component composition, backend graph definition, and distributed execution. To this end, we introduce RLgraph, a library for designing and executing reinforcement learning tasks in both static graph and define-by-run paradigms. The resulting implementations are robust, incrementally testable …

Save Cite **Cited by 13** Related articles All 10 versions

Source: https://scholar.google.com, accessed on 18/10/2022

rlgraph / **rlgraph** Public

RLgraph: Modular computation graphs for deep reinforcement learning

Apache-2.0 license

297 stars 40 forks

Source: https://github.com/rlgraph/rlgraph, accessed on 18/10/2022

UNIVERSITY OF CAMBRIDGE

michaelschaarschmidt … **on 5 Nov 2019**

Source: https://github.com/rlgraph/rlgraph, accessed on 18/10/2022

RLgraph still relevant? 2022…

- Ray RLlib incorporates concepts of RLgraph
  Source: https://docs.ray.io/en/latest/rllib/index.html, accessed 18/10/2022

RLgraph still relevant? 2022…

- Ray RLlib incorporates concepts of RLgraph
  Source: https://docs.ray.io/en/latest/rllib/index.html, accessed 18/10/2022
- Autograph far more capable

UNIVERSITY OF CAMBRIDGE

Critique

- Is being backend agnostic really beneficial?
  - Constant updates with new backend versions necessary
  - Increased maintenance effort

UNIVERSITY OF
CAMBRIDGE

Critique

- Is being backend agnostic really beneficial?
  - Constant updates with new backend versions necessary
  - Increased maintenance effort
- Problems in one specific backend should be addressed in that backend

Critique

- Is being backend agnostic really beneficial?
  - Constant updates with new backend versions necessary
  - Increased maintenance effort
- Problems in one specific backend should be addressed in that backend
- Is mixing Python control flow with machine learning framework code really bad?
  - Autograph

UNIVERSITY OF CAMBRIDGE

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In OSDI, volume 15, pp. 265–283, 2016.
- Chollet, F. et al. Keras. https://keras.io, 2014.
- Facebook Inc. ONNX- Open Neural Network Exchange Format. https://onnx.ai, 2016.

UNIVERSITY OF CAMBRIDGE

References

- Jia, Z., Zaharia, M., and Aiken, A. Beyond data and model parallelism for deep neural networks. http://arxiv.org/pdf/1806.05358v1, 2018.
- Liang, E., Liaw, R., Nishihara, R., Moritz, P., Fox, R., Goldberg, K., Gonzalez, J., Jordan, M., and Stoica, I. Rllib: Abstractions for distributed reinforcement learning. In International Conference on Machine Learning, pp. 3058–3068, 2018.

UNIVERSITY OF CAMBRIDGE

References

- Moritz, P., Nishihara, R., Wang, S., Tumanov, A., Liaw, R., Liang, E., Paul, W., Jordan, M. I., and Stoica, I. Ray: A distributed framework for emerging AI applications. CoRR, abs/1711.05889, 2017. URL http://arxiv.org/abs/1712.05889.
- Palkar, S., Thomas, J. J., Shanbhag, A., Narayanan, D., Pirk, H., Schwarzkopf, M., Amarasinghe, S., Zaharia, M., and InfoLab, S. Weld: A common runtime for high performance data analytics. In Conference on Innovative Data Systems Research (CIDR), 2016.

References

- Plappert, M. keras-rl. https://github.com/matthiasplappert/keras-rl, 2015.
- Schaarschmidt, M., Kuhnle, A., Ellis, B., Fricke, K., Gessert, F., and Yoneki, E. Lift: Reinforcement learning in computer systems by learning from demonstrations. CoRR, abs/1807.07903, 2018. URL http://arxiv.org/abs/1707.06347.
- Sidor, S. and Schulman, J. Openai baselines. website, 2016. URL https://blog.openai.com/ openai-baselines-dqn/.

UNIVERSITY OF CAMBRIDGE

RLgraph: Modular Computation Graphs for Deep Reinforcement Learning

M. Schaarschmidt, S. Mika, K. Fricke, E. Yoneki at SysML, 2019

R244 Large-scale data processing and optimisation
Presentation by Martin Graf on 19/10/2022