Comparison of Hyperparameter Optimization Approaches for Deep Reinforcement Learning

LA-MCTS, SMAC3, Ray-Tune, Gymnasium, CleanRL

Motivation

Deep Learning

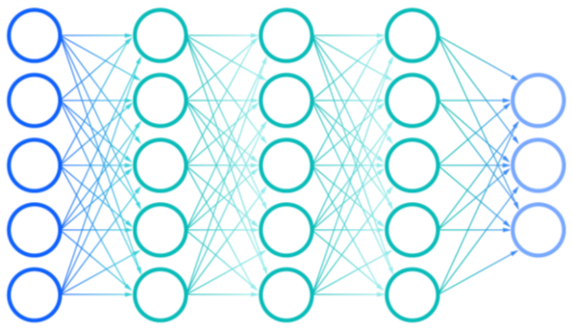

Actor-Critic Networks

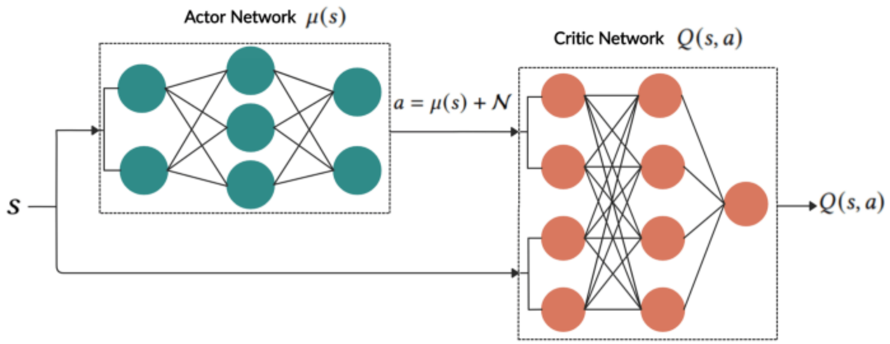

Non-Stationarity

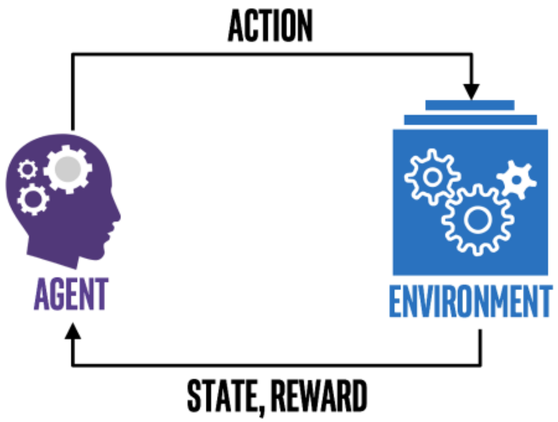

Source: Jaderberg et al., 2017

An active area of research…

Learning Search Space Partition for Black-box Optimization using Monte Carlo Tree Search

Linnan Wang
Brown University
linnan_wang@brown.edu

Rodrigo Fonseca
Brown University
rfonseca@cs.brown

SMAC3: A Versatile Bayesian Optimization Package for Hyperparameter Optimization

Marius Lindauer[1] LINDAUER@TNT.UNI-HANNOVER.DE
GENSPK@CS.UNI-FREIBURG.DE
EURERM@CS.UNI-FREIBURG.DE
EDENKA@CS.UNI-FREIBURG.DE
DENG@TNT.UNI-HANNOVER.DE
AMINS@TNT.UNI-HANNOVER.DE
KOPF@TNT.UNI-HANNOVER.DE
SASS@TNT.UNI-HANNOVER.DE
FH@CS.UNI-FREIBURG.DE
er for Artificial Intelligence

On hyperparameter optimization of machine learning algorithms: Theory and practice

Li Yang, Abdallah Shami

Department of Electrical and Computer Engineering, University of Western Ontario, 1151 Richmond St, London, ON N6A 3K7, Canada

Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization

Lisha Li LISHAL@CS.CMU.EDU
Carnegie Mellon University, Pittsburgh, PA 15213
Kevin Jamieson JAMIESON
University of Washington, Seattle, WA 98195
Giulia DeSalvo
Google Research, New York, NY 10011
Afshin Rostamizadeh
Google Research, New York, NY 10011
Ameet Talwalkar
Carnegie Mellon University, Pittsburgh, PA 15213
Determined AI

Population Based Training of Neural Networks

imon Osindero **Wojciech M. Czarnecki**
Vinyals **Tim Green** **Iain Dunning**
Fernando **Koray Kavukcuoglu**
London, UK

Hyper-Parameter Optimization: A Review of Algorithms and Applications

Tong Yu YUTONG01@INSPUR.COM
Department of AI and HPC
Inspur Electronic Information Industry Co., Ltd
1036 Langchao Rd, Jinan, Shandong, China

Hong Zhu ZHUHONGBJ@INSPUR.COM
Department of AI and HPC
Inspur (Beijing) Electronic Information Industry Co., Ltd
2F, Block C, 2 Xinxi Rd., Shangdi, Haidian Dist, Beijing, China

Open-Source Project

Approaches

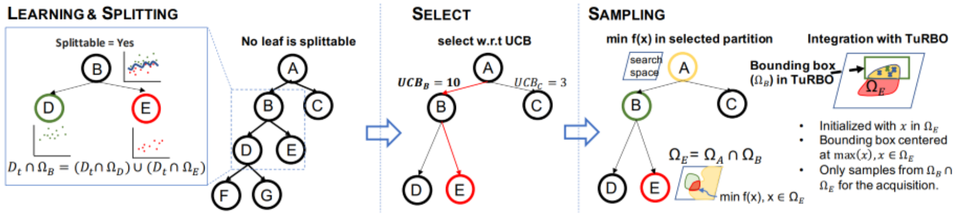

LA-MCTS

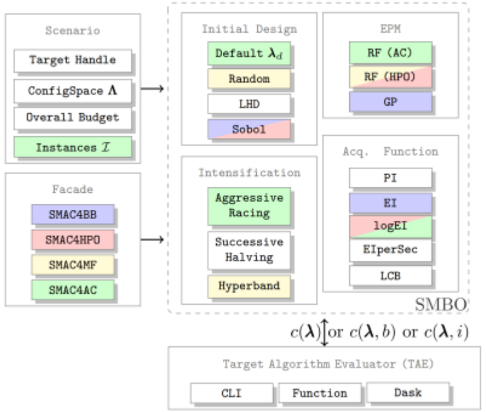

SMAC3

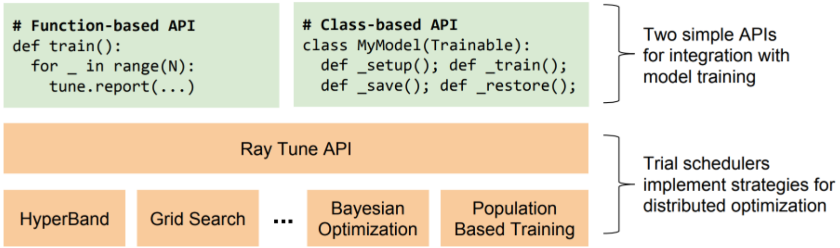

Ray-Tune

Environments: Gymnasium

Gymnasium

RL Algorithms: CleanRL

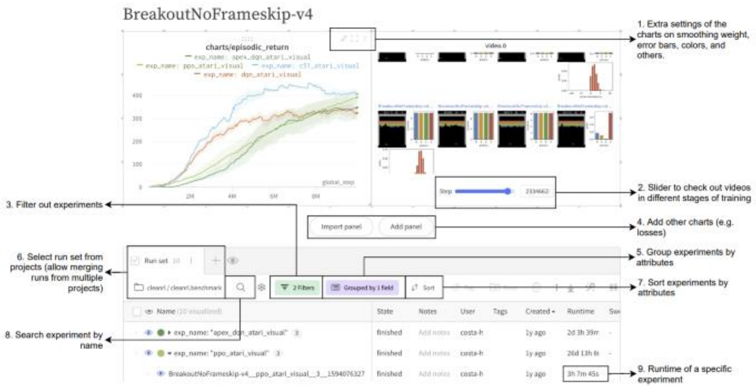

Latent Action Monte Carlo Tree Search (LA-MCTS)

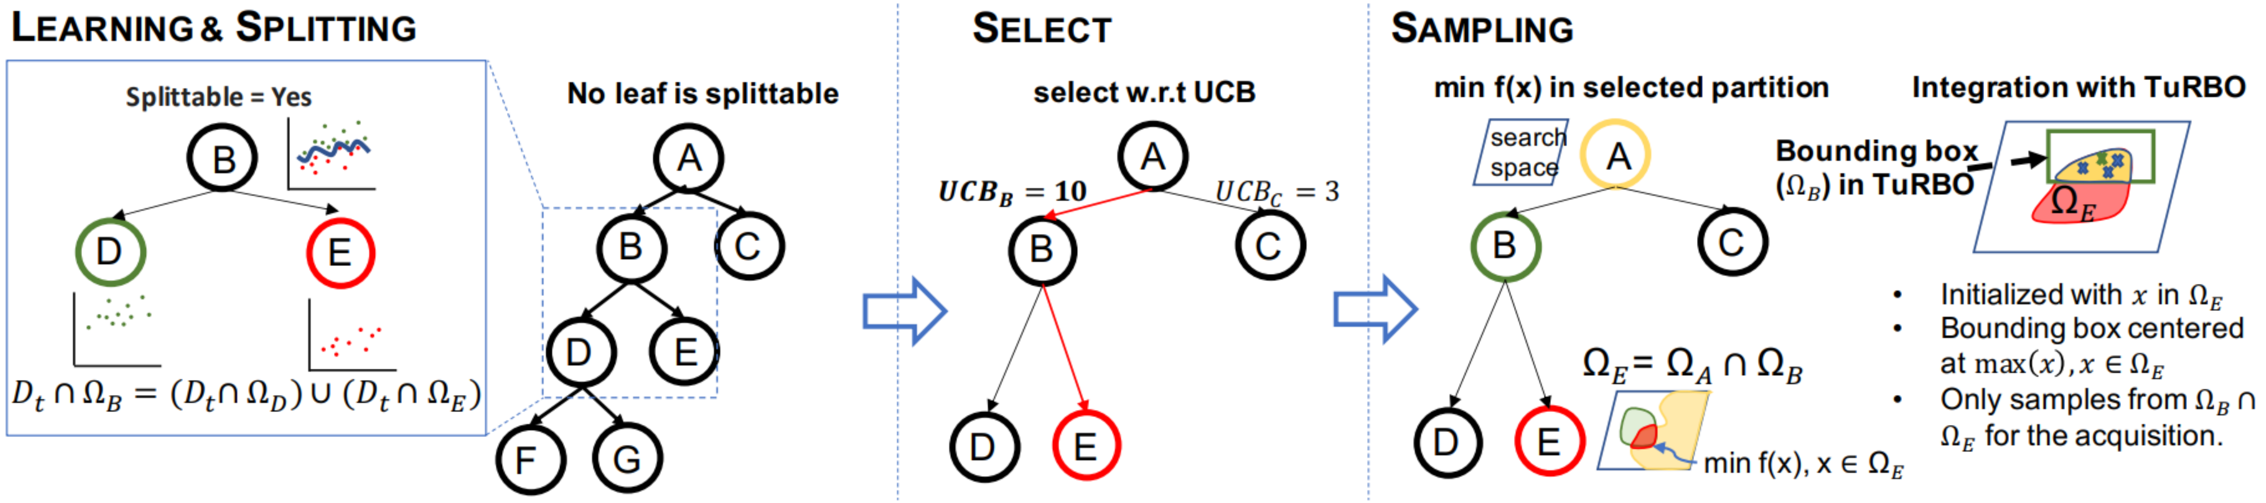

Source: Wang et al., 2020

Sequential Model Algorithm Configuration (SMAC3)

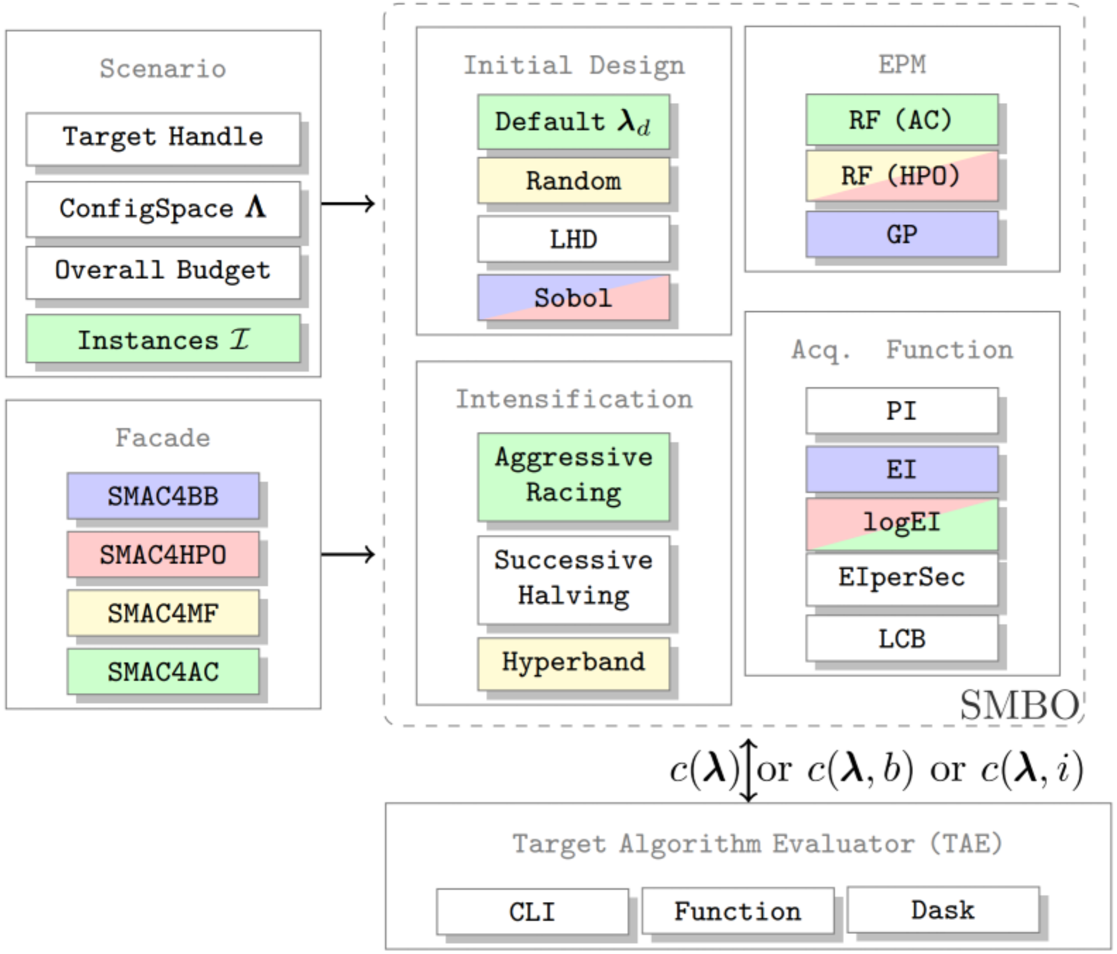

Source: Lindauer et al., 2022

Ray-Tune

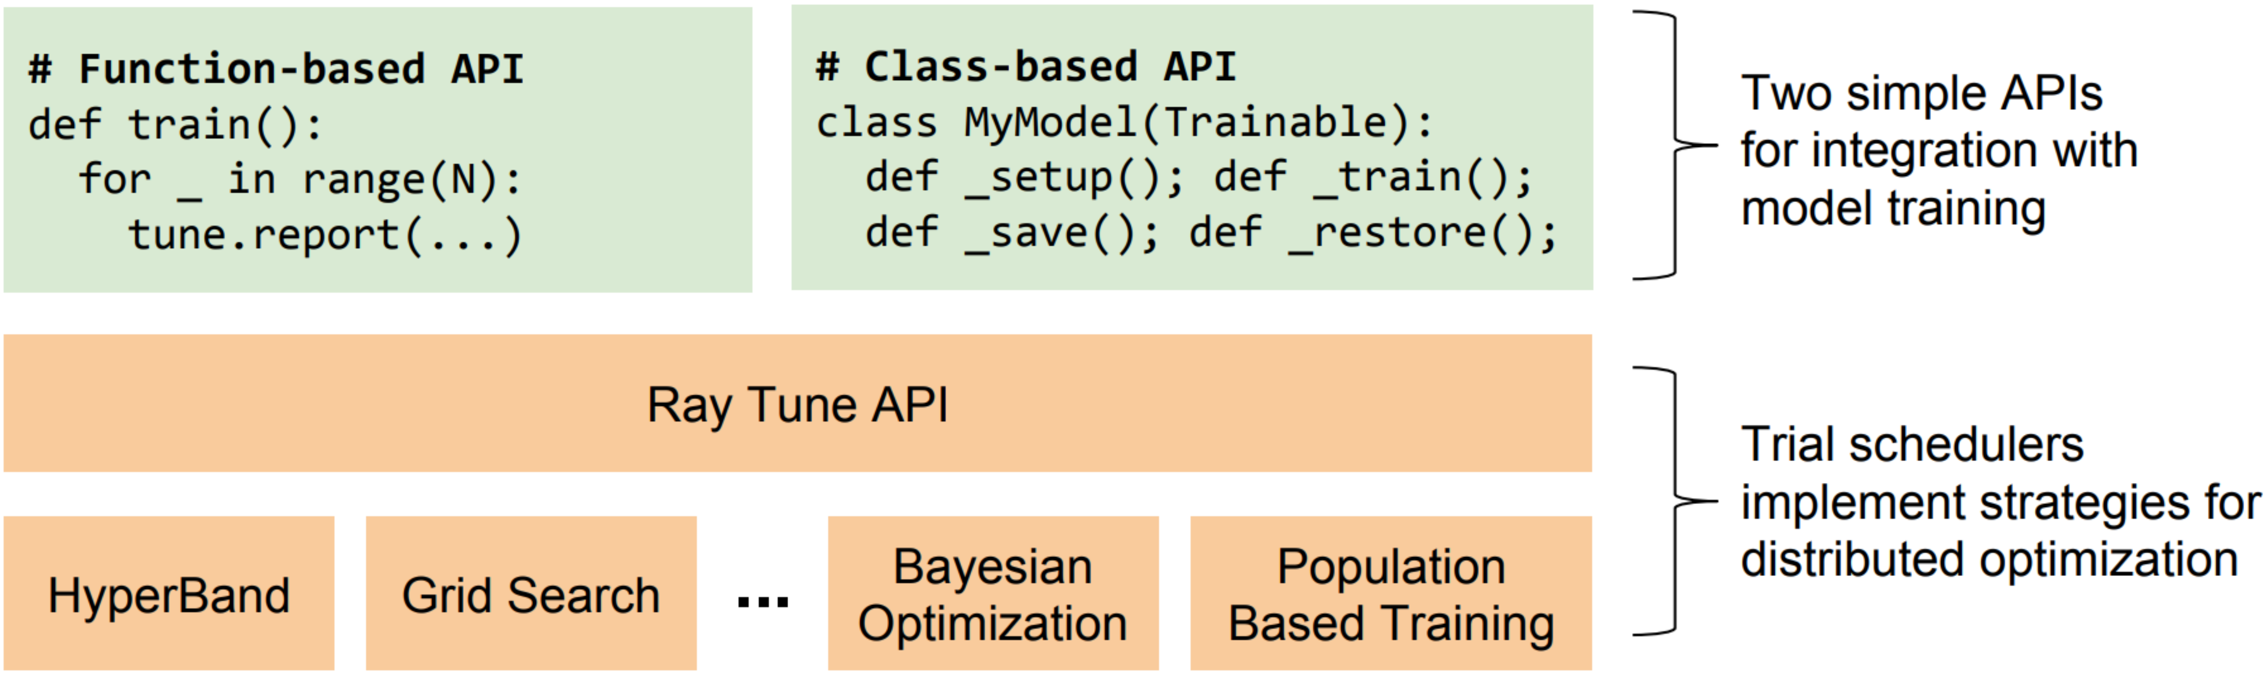

Source: Liaw et al., 2018

Comparison

Environments

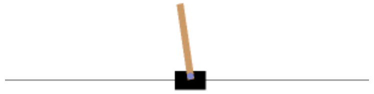

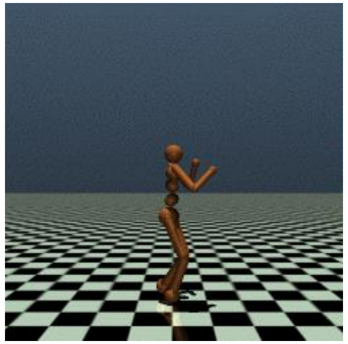

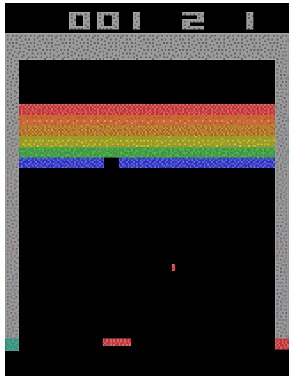

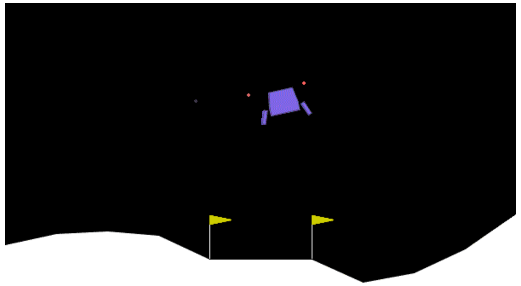

Algorithms

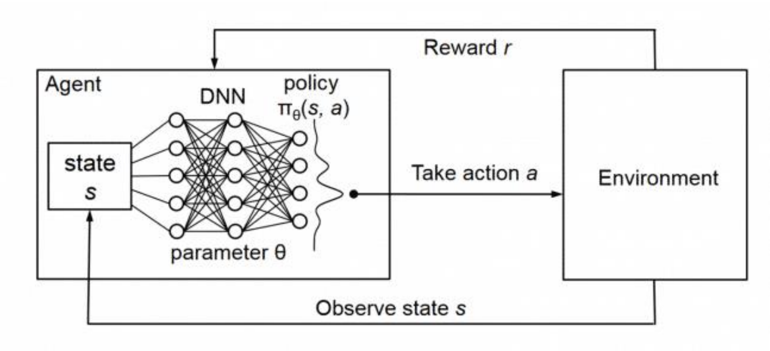

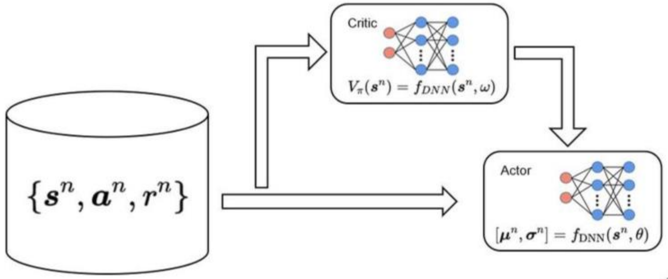

Scalability

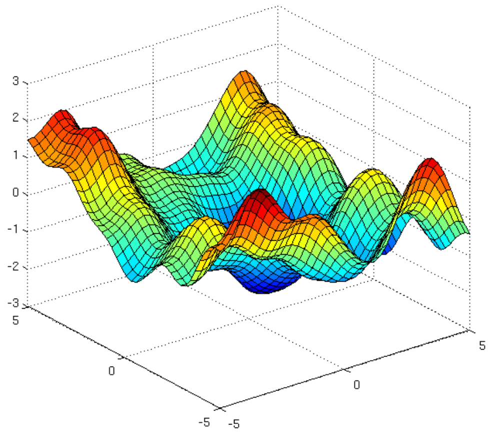

Work Plan

- Experimental design and setup
- Implementation of common optimizer interface
- Run experiments
  - LA-MCTS
  - SMAC3
  - Ray-Tune
- Evaluation on distributed setup
- If time permits, extension to the Meta-World benchmark

Questions / Discussion

References

- L. Wang, R. Fonseca, and Y. Tian, “Learning search space partition for black-box optimization using monte carlo tree search,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, Dezember 2020, pp. 19511–19522.
- M. Lindauer *et al.*, “SMAC3: A Versatile Bayesian Optimization Package for Hyperparameter Optimization,” *Journal of Machine Learning Research*, vol. 23, no. 54, pp. 1–9, 2022.
- R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica, “Tune: A Research Platform for Distributed Model Selection and Training.” arXiv, Jul. 13, 2018. doi: 10.48550/arXiv.1807.05118.
- L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, “Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization,” *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 6765–6816, Jan. 2017.
- M. Jaderberg *et al.*, “Population Based Training of Neural Networks.” arXiv, Nov. 28, 2017. doi: 10.48550/arXiv.1711.09846.
- L. Yang and A. Shami, “On hyperparameter optimization of machine learning algorithms: Theory and practice,” *Neurocomputing*, vol. 415, pp. 295–316, 2020, doi: https://doi.org/10.1016/j.neucom.2020.07.061.
- T. Yu and H. Zhu, “Hyper-Parameter Optimization: A Review of Algorithms and Applications.” arXiv, Mar. 12, 2020. doi: 10.48550/arxiv.2003.05689.
- G. Brockman *et al.*, “OpenAI Gym,” Jun. 2016, doi: 10.48550/arXiv.1606.01540.
- S. Huang *et al.*, “CleanRL: High-quality Single-file Implementations of Deep Reinforcement Learning Algorithms,” *Journal of Machine Learning Research*, vol. 23, no. 274, pp. 1–18, 2022.
- T. Yu *et al.*, “Meta-World: A Benchmark and Evaluation for Multi-Task and Meta Reinforcement Learning,” in *Proceedings of the Conference on Robot Learning*, Nov. 2020, vol. 100, pp. 1094–1100. [Online]. Available: https://proceedings.mlr.press/v100/yu20a.html

Image Sources

- L. Wang, R. Fonseca, and Y. Tian, “Learning search space partition for black-box optimization using monte carlo tree search,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, Dezember 2020, pp. 19511–19522.
- M. Lindauer *et al.*, “SMAC3: A Versatile Bayesian Optimization Package for Hyperparameter Optimization,” *Journal of Machine Learning Research*, vol. 23, no. 54, pp. 1–9, 2022.
- R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica, “Tune: A Research Platform for Distributed Model Selection and Training.” arXiv, Jul. 13, 2018. doi: 10.48550/arXiv.1807.05118.
- L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, “Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization,” *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 6765–6816, Jan. 2017.
- M. Jaderberg *et al.*, “Population Based Training of Neural Networks.” arXiv, Nov. 28, 2017. doi: 10.48550/arXiv.1711.09846.
- L. Yang and A. Shami, “On hyperparameter optimization of machine learning algorithms: Theory and practice,” *Neurocomputing*, vol. 415, pp. 295–316, 2020, doi: https://doi.org/10.1016/j.neucom.2020.07.061.
- T. Yu and H. Zhu, “Hyper-Parameter Optimization: A Review of Algorithms and Applications.” arXiv, Mar. 12, 2020. doi: 10.48550/arxiv.2003.05689.
- S. Huang *et al.*, “CleanRL: High-quality Single-file Implementations of Deep Reinforcement Learning Algorithms,” *Journal of Machine Learning Research*, vol. 23, no. 274, pp. 1–18, 2022.
- https://1.cms.s81c.com/sites/default/files/2021-01-06/ICLH_Diagram_Batch_01_03-DeepNeuralNetwork-WHITEBG.png
- https://i0.wp.com/upaspro.com/wp-content/uploads/2020/08/1y_oM_CnrT2_-7hTlq_pe_g.png?w=1400&ssl=1
- https://miro.medium.com/max/1400/0*WC4l7u90TsKs_eXj.png
- https://raw.githubusercontent.com/Farama-Foundation/Gymnasium/main/gymnasium-text.png
- https://www.gymlibrary.dev/_images/breakout.gif
- https://www.gymlibrary.dev/_images/cart_pole.gif
- https://www.gymlibrary.dev/_images/humanoid.gif
- https://www.gymlibrary.dev/_images/lunar_lander.gif
- https://www.novatec-gmbh.de/wp-content/uploads/reinforcement_learning_loop-650x294.png
- https://www.researchgate.net/publication/359917910/figure/fig3/AS:1144311030788098@1649836418374/The-proximal-policy-optimization-in-domain-control-scheme-Algorithm-1-Proximal-Policy.jpg
- https://thegradient.pub/content/images/size/w1600/2019/11/kernel_cookbook-2.png
- https://images.squarespace-cdn.com/content/v1/596f25c2725e25fb89b3a6f4/1544176431098-WT0ELR52Q0X1IQC89YRH/discussion+cc3.0.png?format=1500w