

## MIT Open Access Articles

*Rlpy: A Value-Function-Based Reinforcement Learning Framework for Education and Research*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Geramifard, Alborz et al. "RLPy: A Value-Function-Based Reinforcement Learning Framework for Education and Research." *Journal of Machine Learning Research* 16 (2015):1573–1578.

**As Published:** <http://jmlr.org/papers/v16/geramifard15a.html>

**Publisher:** MIT Press

**Persistent URL:** <http://hdl.handle.net/1721.1/105742>

**Version:** Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

**Terms of Use:** Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



# RLPy: A Value-Function-Based Reinforcement Learning Framework for Education and Research

**Alborz Geramifard**<sup>12</sup>

AGF@CSAIL.MIT.EDU

*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology,  
77 Massachusetts Ave., Cambridge, MA 02139 – USA*

**Christoph Dann**<sup>1</sup>

CDANN@CMU.EDU

*Machine Learning Department, Carnegie Mellon University,  
5000 Forbes Ave., Pittsburgh, PA 15213 – USA*

**Robert H. Klein**<sup>1</sup>

BOBKLEIN2@ALUM.MIT.EDU

*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology,  
77 Massachusetts Ave., Cambridge, MA 02139 – USA*

**William Dabney**<sup>2</sup>

WDDABNEY@AMAZON.COM

*Amazon.com,*

*440 Terry Ave. N, Seattle, WA 98109 – USA*

**Jonathan P. How**

JHOW@MIT.EDU

*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology,  
77 Massachusetts Ave., Cambridge, MA 02139 – USA*

**Editor:** Geoff Holmes

## Abstract

RLPy is an object-oriented reinforcement learning software package with a focus on value-function-based methods using linear function approximation and discrete actions. The framework was designed for both educational and research purposes. It provides a rich library of fine-grained, easily exchangeable components for learning agents (e.g., policies or representations of value functions), facilitating recently increased specialization in reinforcement learning. RLPy is written in Python to allow fast prototyping, but is also suitable for large-scale experiments through its built-in support for optimized numerical libraries and parallelization. Code profiling, domain visualizations, and data analysis are integrated in a self-contained package available under the Modified BSD License at <http://github.com/rlpy/rlpy>. All of these properties allow users to compare various reinforcement learning algorithms with little effort.

**Keywords:** reinforcement learning, value-function, empirical evaluation, open source

## 1. Introduction

An integral part of most artificial intelligence courses are value-function-based methods using linear function approximation for solving Markov decision processes (such as linear Q-learning or SARSA). In addition, many researchers build upon this well-understood and

- 
1. The first three authors contributed equally to this work.
  2. The majority of this work was done prior to Amazon involvement of the authors. This paper does not reflect the views of the Amazon company.

powerful framework and aim at improving existing methods by, for example, feature learning (Keller et al., 2006; Parr et al., 2007; Geramifard et al., 2011), or policies with better exploration-exploitation trade-off (Nouri and Littman, 2009; Jaksch et al., 2010; Li, 2012).

The need to unify and increase reusability of software packages for reinforcement learning research has been widely discussed (Tanner and White, 2009; Schaul et al., 2010), and many successful tools have been created (see Section 2). However, it is desirable to have a software framework that is 1) easily accessible by novices so that they may compare and understand existing algorithms, and 2) efficient for researchers who perform large scale experiments and advance the state-of-the-art.

By focusing on the prominent class of value-function-based methods with linear function approximation using discrete actions, RLPy aims at being such a software framework that provides simple and convenient tools for conducting sequential decision making experiments. In the following, we present the main features of RLPy and highlight those that distinguish it from existing frameworks.

## 2. Existing Frameworks

The following existing software packages have some overlap with RLPy:

1. *RL-Toolbox*: (Neumann, 2005) C++ RL toolbox focusing on continuous state-spaces
2. *CLSsquare*: (Riedmiller et al., 2012) C++ RL framework focused on interfaces with several robotics
3. *libPG*: (Aberdeen, 2007) RL library focused on high-performance policy-gradient algorithm implementations
4. *rllib* (Frezza-Buet and Geist, 2013): Template-based C++ RL library for value-function methods
5. *rl-explore-ros-pkg*: (Hester, 2013) ROS package for RL algorithms
6. *JRLF*: (Kochenderfer, 2006) Small-scale Java-Framework for RL experiments
7. *PIQLE*: (de Comite, 2006) Java-Framework for RL experiments
8. *RLPark*: (Degris, 2013) Java reinforcement learning library
9. *RLLib*: (Abeyruwan, 2013) Port of RLPark into C++
10. *RL-Glue, RL-Library*: (Tanner and White, 2009) Protocol for RL experiments and reference implementations
11. *ApproxRL*: (Busoniu, 2010) Matlab Toolbox with RL and dynamic programming algorithms
12. *MMLF*: (Metzen and Edgington, 2011) Python-based framework for reinforcement learning
13. *PyBrain* (Schaul et al., 2010): Machine learning library focused on neural networks with RL support

For the sake of brevity, we do not compare RLPy against each of the existing frameworks in detail but highlight key differences in the following section by referencing the list above.

## 3. Why RLPy?

*Improved Granularity of Agents with Linear Value Functions.* RL has advanced significantly over the past decade, leading researchers to narrow their focus towards specialized, independent aspects of RL agents, such as approximate function representations, exploration schemes, and learning rates. The structure of numerous existing frameworks (2, 6, 7, 10, 12) does not properly account for this increased specialization and makes it cumbersome to exchange, for example, the way the value function is represented in a learning agent. RLPy addresses this issue by separating these components into exchangeable classes (shown as green boxes in Figure 1) and other minor components such as learning rates into separate functions. This division reduces implementation effort, promotes reusability, and facilitates automated testing. Code for an example experiment that exploits this modularity is shown in Figure 2. In addition, the assumption of linearly parameterizing the value function allows RLPy to provide many tools and helpers for designing state features. For example, in large




Figure 1: RLPy framework - Green components constitute an RL agent which did not exist as separate components in previous RL frameworks. The experiment module handles the interaction between the agent and the domain; gray arrows depict the information flow in a conventional RL framework (Sutton and Barto, 1998).

```
import rlp
#### Domain ####
domain = rlp.Domains.InfCartPoleBalance()
### Agent ###
representation = rlp.Representations.Tabular(domain, discretization=20)
policy = rlp.Policies.eGreedy(representation, epsilon=0.1)
agent = rlp.Agents.SARSA(policy, representation, domain.discount_factor)
### Experiment ####
experiment = rlp.Experiments.Experiment(agent, domain, max_steps=100000)
experiment.run()
experiment.save()
```

Figure 2: RLPy code for setting up and running an experiment: SARSA learning for 100,000 steps how to balance an inverted pole on a cart while following an  $\epsilon$ -greedy policy and using discretized tabular features.

MDPs where using a tabular representation is infeasible, the `IndependentDiscretization` representation creates features by ignoring dependency among dimensions of states.

*Rapid Prototyping with Python.* RLPy is fully object-oriented and based primarily on the Python language (van Rossum and de Boer, 1991). Low-level, computationally-intensive tools are implemented in Cython (a compiled and typed version of Python) or C++. In contrast to other packages (1 – 9) written solely in C++ or Java, this approach leverages the user-friendliness, conciseness, and portability of Python while supplying computational efficiency where needed. This combination allows researchers to prototype new ideas quickly and comfortably without sacrificing the computing speed necessary to conduct large-scale experiments. In addition, the Python-based approach of RLPy is particularly suited for education as it does not require any proprietary software (in contrast to 11).

*“Batteries Included” – Many Existing Components and Benchmarks.* RLPy includes an ever-growing repository of components which may be combined to form new RL agents. While many frameworks (1, 3, 4, 6) only include classic benchmark domains such as *Puddle-World* or an *Inverted Pendulum on Cart*, RLPy supplies a large number of more challenging domains such as *HIV-Treatment*, *Hovering a Helicopter*, and *Pac-Man*. In addition to implementations of most value-function-based RL algorithms, RLPy includes experimental support for dynamic programming methods that require full domain knowledge but yield optimal policies. This is especially useful as a baseline for comparison with (often sub-optimal) policies generated by RL agents.




Figure 3: RLPy sample outputs of RLPy plotting (left) and profiling (right) tools: A portion of the profiling graph of the example code (Figure 2) in which the green box shows the statistics of executing the `learn` function  $10^5$  times. It required 37.38% of the CPU-time for completion, out of which its main body was responsible only for 7.15% of the computation while the rest was spent in other called functions.

*Ease of Use and Development.* Numerous tools are shipped with RLPy that facilitate ease of use and efficiency. One example is the code profiler, which produces a visual runtime graph of the source code (c.f. Figure 3 right) and identifies slow routines. This information allows the researcher to reduce the runtime of an algorithm with minimal effort and discourages premature runtime optimization. Additionally, every RLPy domain has a visualization, an important feature lacking in other frameworks (3, 4, 7). These visuals help user quickly assess and gain intuition about the algorithm and domain behavior.

*Automation of Experiments.* RLPy aims to promote reproducible research. To this end, it provides a suite of tools to automate the entire experiment pipeline. For example, RLPy allows concise specification of experiment settings (see Figure 2) and automated and efficient hyperparameter optimization with the `hyperopt` package (Yamins et al., 2013). Researchers can share their experimental setups by publishing short settings files, and colleagues can reproduce the results when running the scripts independent of their hardware or operating system. Additionally, RLPy experiments are natively parallelizable. Once parameters are selected, the user simply specifies the number of CPU cores RLPy can utilize for multiple experiments to test statistical significance. RLPy enables further scaling by switching seamlessly from a single machine to a job-based cluster (e.g. HTCCondor) while ensuring results remain identical across varying hardware. RLPy also provides automated tools for generation of final publication-ready plots of results (see Figure 3 left); researchers need only specify the quantities that should appear on the plot. To the best of our knowledge this degree of automation of the entire experimentation pipeline is unique to RLPy.

## 4. Conclusion

RLPy is a new reinforcement learning framework focused on value-function-based reinforcement learning using linear function approximation with discrete actions. It simplifies the construction of learning agents and makes it easier for novices and experts alike to evaluate and compare algorithms, representations, environments, and other RL components. RLPy also provides many tools for conducting reproducible experiments from initial prototyping to final plotting. The framework is entirely open-source and all contributions are welcome.

## References

- Douglas Aberdeen. LibPGRL: A high performance reinforcement learning library in C++, 2007. URL <https://code.google.com/p/libpgrl>.
- Saminda Abeyruwan. RLLib reinforcement learning c++ template library, 2013. URL <http://web.cs.miami.edu/home/saminda/rllib.html>.
- Lucian Busoniu. Approxrl: A matlab toolbox for approximate reinforcement learning and dynamic programming, 2010. URL [http://busoniu.net/files/repository/readme\\_approxrl.html](http://busoniu.net/files/repository/readme_approxrl.html).
- Francesco de Comitè. PIQLE: A platform for implementation of Q-learning experiments, 2006. URL <http://piqle.sourceforge.net>.
- Thomas Degris. RLPark, 2013. URL <http://rlpark.github.io>.
- Herve Frezza-Buet and Matthieu Geist. A C++ template-based reinforcement learning library: Fitting the code to the mathematics. *Journal of Machine Learning Research (JMLR)*, 14:625–628, 2013.
- Alborz Geramifard, Finale Doshi, Joshua Redding, Nicholas Roy, and Jonathan How. Online discovery of feature dependencies. In *International Conference on Machine Learning (ICML)*, 2011.
- Todd Hester. rl-texplore-ros-pkg: Reinforcement learning framework, agents, and environments with ROS interface, 2013. URL <https://code.google.com/p/rl-texplore-ros-pkg>.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research (JMLR)*, 11:1563–1600, 2010.
- Philipp W. Keller, Shie Mannor, and Doina Precup. Automatic basis function construction for approximate dynamic programming and reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2006.
- Mykel Kochenderfer. JRLF: Java reinforcement learning framework, 2006. URL <http://mykel.kochenderfer.com/jrlf>.
- Lihong Li. Sample complexity bounds of exploration. In Marco Wiering and Martijn van Otterlo, editors, *Reinforcement Learning: State of the Art*. Springer Verlag, 2012.
- Jan Hendrik Metzen and Mark Edgington. Maja machine learning framework, 2011. URL <http://mmlf.sourceforge.net>.
- Gerhard Neumann. The reinforcement learning toolbox, reinforcement learning for optimal control tasks, 2005.
- Ali Nouri and Michael L. Littman. Multi-resolution exploration in continuous spaces. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.

- Ronald Parr, Christopher Painter-Wakefield, Lihong Li, and Michael Littman. Analyzing Feature Generation for Value-Function Approximation. In *International Conference on Machine Learning (ICML)*, 2007.
- Martin Riedmiller, Manuel Blum, and Thomas Lampe. CLS<sup>2</sup>: Closed loop simulation system, 2012. URL <http://ml.informatik.uni-freiburg.de/research/clsquare>.
- Tom Schaul, Justin Bayer, Daan Wierstra, Yi Shun, Martin Felder, Frank Sehnke, Thomas Rückstieß, and Jürgen Schmidhuber. PyBrain. *Journal of Machine Learning Research (JMLR)*, 11:743–746, 2010.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
- Brian Tanner and Adam White. RL-Glue : Language-independent software for reinforcement-learning experiments. *Journal of Machine Learning Research (JMLR)*, 10:2133–2136, 2009.
- Guido van Rossum and Jelke de Boer. Interactively testing remote servers using the python programming language. *CWI Quarterly*, 4(4):283–303, 1991.
- Daniel Yamins, David Tax, and James S. Bergstra. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International Conference on Machine Learning (ICML)*, 2013.