

REINFORCEMENT LEARNING RECORDBOOK < RL @ HOME >

The reinforcement learning recordbook ([RL-Recordbook](#)) is an open source project that facilitates empirical comparison in the reinforcement learning community. The RL-Recordbook provides an experimental framework that allows painless, faithful reproductions of experimental results. This framework can reduce the bias and effort often associated with re-implementing and tuning alternative algorithms for empirical comparison in reinforcement learning. The framework (illustrated in the figure) is implemented as a set of loosely-coupled modules that can be run distributed across the Internet, within a local network, or on a single machine.

Our experimental framework is augmented by a large public repository of results and source code for reinforcement learning agents, environments, and experiments. This repository will be useful when planning future empirical comparisons: instead of designing a new experiment, a researcher may find an established experiment that measures the relevant criteria for evaluating their new idea. The researcher can run their algorithm with this established experiment and simply download the results for the appropriate comparison algorithms. This removes the bias that creeps into experiments because of incorrect implementation or configuration of the alternative algorithm, allowing the researcher to focus on the correctness and performance of their own algorithm. It also saves time and improves understanding because the source code for the experiment and alternate algorithms can be published for future extension and re-use.

The RL-Recordbook may grow to consist of millions of results, spanning thousands of parameter combinations of hundreds of algorithms in various experiments with different reinforcement learning environments. A project of this magnitude has computation, storage, and availability implications that scale beyond what can reliably be hosted with a traditional web service/site/application hosted on university servers or affordably with traditional hosting services. Instead, we've designed the system to be completely modular and distributed, with various components communicating via distributed queue, database, and file services. One of our implementations of the recordbook uses [Amazon's](#) distributed storage ([S3](#)), queue ([SQS](#)), and database ([SimpleDB](#)) services. This implementation is called RL@Home, after the [SETI@Home](#) project. With RL@Home, the submission, computation, aggregation, storage, and display of the experiments is all performed by loosely-coupled modules. This includes an accounting system so that a user can earn credits for running experiments, and can later spend those credits to have other users run their experiments. The RL-Recordbook is layered on top of other open source projects that have been spearheaded by our group, including [RL-Glue](#), [RL-Viz](#), and [RL-Library](#), featured at a [NIPS demo](#), and at the [MLOSS workshop](#).

Our poster at the [parallel implementations workshop](#) will introduce the community to the first beta-release of the RL-Recordbook and RL@Home. The poster is meant to raise awareness and interest in the project, so we will also focus on the new research opportunities to do macro-analysis of reinforcement learning agent performance across problems and parameter combinations. We will also discuss the opportunity to create macro-algorithms that improve the database by predicting interesting experimental results and autonomously submitting those experiments into the work queue. We will present new macro-insights gained from the first large scale study done with the RL-Recordbook: running tens of thousands of parameter and function approximation combinations of the Sarsa(λ) algorithm on several standard benchmark reinforcement learning problems.

RL@Home Data and Work Flow

