

Paper Submission

Authors are encouraged to submit high-quality, original work that has neither appeared in, nor is under consideration by, other journals.


Springer offers authors, editors and reviewers of **Machine Learning** a web-enabled online manuscript submission and review system. Our online system offers authors the ability to track the review process of their manuscript.

Manuscripts should be submitted to: <http://MACH.edmgr.com>. This online system offers easy and straightforward log-in and submission procedures, and supports a wide range of submission file formats. When submitting please be sure to choose the manuscript type, "SI: Empirical RL."

Important Dates

- Paper submission deadline: February 26, 2010
- Notification of acceptance: June 30, 2010
- Final manuscript: September 30, 2010

www.Springer.com/10994



ISSN: 0885-6125 (print)
1573-0565 (electronic)

Editor-in-Chief:

Foster Provost,
New York University

Special Issue Call for Papers

Empirical Evaluations in Reinforcement Learning

Shimon Whiteson
Informatics Institute
University of Amsterdam
s.a.whiteson@uva.nl

Michael Littman
Department of Computer Science
Rutgers University
mlittman@cs.rutgers.edu

The continuing development of a field requires a healthy exchange between theoretical advances and experimental observations. The purpose of this special issue is to assess progress in empirical evaluations of reinforcement-learning algorithms and to encourage the adoption of effective experimental methodologies. The last several years have seen new trends in uniform software interfaces between environments and learning algorithms, community comparisons and competitions, and an increased interest in experimenting with reinforcement learning in embedded systems. We enthusiastically solicit papers on relevant topics such as:

- The design and dissemination of standardized frameworks and repositories for algorithms, methods, and/or results.
- Experience of organizers and participants in reinforcement-learning competitions and bake-offs.
- Novel evaluation methodologies or metrics.
- Careful empirical comparisons of existing methods.
- Novel methods validated with strong empirical results on existing benchmarks, especially those used in recent RL Competitions (see <http://www.rl-competition.org/>).
- Applications of reinforcement-learning approaches to real-life environments such as computer networks, system management and robotics.
- Theoretical work such as sample complexity bounds that can be used to guide the design of benchmarks and evaluations.

The emphasis of the special issue is not on the development of novel algorithms. Instead, papers will be assessed in terms of the insights they provide about how best to assess performance in reinforcement learning, i.e., the "meta" problem of evaluating the evaluation methodologies themselves. In particular, papers presenting empirical results should also discuss what those results reveal about the strengths and weaknesses of the evaluation methodology. Similarly, papers describing real-life applications should make clear what limitations the application exposes in 'off-the-shelf' methods, how the employed method had to be modified to address real-world complications, and what the results show that could not be learned from experiments in 'toy' domains. Papers proposing new evaluation methodologies should include illustrative empirical results offering insights that would be difficult to obtain with conventional methodologies. Finally, papers proposing new evaluation methodologies should also compare and contrast with methodologies in related areas, e.g. supervised learning, explaining why such methodologies are not adequate and what ideas, if any, can be borrowed from them.