

OpenDP a free Reinforcement Learning toolbox for discrete time control problems

OpenDP (<http://opendp.sourceforge.net>) is a free software (under GPL) written in C++ for control problems: (i) with discrete time steps (either finite horizon, or approximation of infinite horizon by temporal-ring for stationary problems), (ii) where the transition function $newState = T(state, decision)$ can be encoded in C++, (iii) possibly stochastic (with the function T above depending on some random process).

OpenDP is designed to be a general framework for all control problems of the previous form. One can easily add new solving methods to the framework and can compare the results to already existing methods. OpenDP includes some benchmark problems, with parameterizable dimension and difficulty. Some benchmarks are stochastic and some are deterministic. Other benchmark problems can easily be added.

OpenDP includes policy search methods in their simplest way, but its core solving method is stochastic dynamic programming. In this method, the Bellman value is computed backward in time by iterating: (i) sampling points in the state domain, (ii) for each point, compute the best expected value given the Bellman function, (iii) the examples and values given (i) and (ii) to learn a regression which gives the next Bellman function.

For each of the three steps considered above, a toolbox is provided including classical methods. For the sampling step, we provide classes generating some classical quasi random sequences (low discrepancy sequences with or without scrambling, low dispersion, ...). For the continuous-optimisation steps, OpenDP is interfaced with famous optimisation frameworks as Opt++, Coin DFO, EO, Beagle, which include for example quasi-Newton algorithms, Hooke&Jeeves, CMAES, SAES and some other algorithms. For the learning step, OpenDP is interfaced with the Weka library and so we include all the regressions implemented in Weka. OpenDP is also interfaced with the Torch library (including SVM and Neural Networks) and the FANN library. The algorithms included in the interfaced toolboxes are available as plugins and so OpenDP can be used even if the libraries are not present.

The download area of OpenDP can be found at <http://opendp.sourceforge.net/>. OpenDP is in beta stage of development. Some typical command-lines provide extensive experiments (see the webpage above), providing a simple way of testing modifications of the source code. OpenDP also comes with a graphical user interface written in Qt, with curve visualizations in 2D, and Bellman function visualization in 3D.

The OpenDP source code is cross-platform. OpenDP actually works under Linux and MS Windows and binary-packages can be downloaded for these platforms. The building process and some mechanisms have to be ported to Mac OS X.

Acknowledgment

Authors are Sylvain Gelly¹, Olivier Teytaud¹, Jérémie Mary². This work was supported in part by the Pascal Network of Excellence.

¹TAO (INRIA), LRI, UMR CNRS, Univ. Paris-Sud; email firstname.lastname@lri.fr.

²Grappa, Univ. Lille UMR CNRS, Jeremie.Mary@grappa.univ-lille3.fr