

mlrMBO: A Toolbox for Model-Based Optimization of Expensive Black-Box Functions

Jakob Richter

TU Dortmund University

Abstract: Many practical optimization tasks, such as finding best parameters for simulators in engineering or hyperparameter optimization in machine learning, are of a black-box nature, i.e., neither formulas of the objective nor derivative information is available. Instead, we can only query the box for its objective value at a given point. If such a query is very time-consuming, the optimization task becomes extremely challenging, as we have to operate under a severely constrained budget of function evaluations. A modern approach is sequential model based-optimization, aka Bayesian optimization. Here, a surrogate regression model learns the relationship between decision variables and objective outcome. Sequential point evaluations are planned to simultaneously exploit the so far learnt functional landscape and to ensure exploration of the search space. A popular instance of this general principle is the EGO algorithm, which uses Gaussian processes coupled with the expected improvement criterion for point proposal. The mlrMBO package offers a rich interface to many variants of model-based optimization. As it builds upon the mlr package for machine learning in R, arbitrary surrogate regression models can be applied. It offers a wide variety of options to tackle different black-box scenarios:

- Optimization of pure continuous as well as mixed continuous-categorical search spaces.
- Single criteria optimization or approximated Pareto fronts for multi-criteria problems.
- Single point proposal or parallel batch point planning during optimization.

The package is designed as a convenient, easy-to-use toolbox of popular state-of-the-art algorithms, but can also be used as as a research framework for algorithm designers.

Keywords: machine learning, hyperparameter optimization, tuning, black box optimization, bayesian optimization